

A framework for the local information dynamics of distributed computation in complex systems

Joseph T. Lizier,^{1,2,*} Mikhail Prokopenko,¹ and Albert Y. Zomaya²

¹*CSIRO Information and Communications Technology Centre,
Locked Bag 17, North Ryde, NSW 1670, Australia*

²*School of Information Technologies, The University of Sydney, NSW 2006, Australia
(Dated: November 17, 2008)*

The nature of distributed computation has often been described in terms of the component operations of universal computation: *information storage*, *transfer* and *modification*. We introduce the first complete framework that quantifies each of these individual information dynamics on a local scale within a system, and describes the manner in which they interact to create non-trivial computation where “the whole is greater than the sum of the parts”. We apply the framework to cellular automata, a simple yet powerful model of distributed computation. In this application, the framework is demonstrated to be the first to provide quantitative evidence for several important conjectures about distributed computation in cellular automata: that blinkers embody information storage, particles are information transfer agents, and particle collisions are information modification events. The framework is also used to investigate and contrast the computations conducted by several well-known cellular automata, highlighting the importance of information coherence in complex computation. Our results provide important quantitative insights into the fundamental nature of distributed computation and the dynamics of complex systems, as well as impetus for the framework to be applied to the analysis and design of other systems.

PACS numbers: 89.75.Fb, 89.75.Kd, 89.70.Cf, 05.65.+b

Keywords: distributed computation, information storage, information transfer, information modification, information theory, complex systems, self-organization, cellular automata, particles, gliders, domains

I. INTRODUCTION

The nature of distributed computation has long been a topic of interest in complex systems science, physics, artificial life and bioinformatics. In particular, emergent complex behavior has often been described from the perspective of computation within the system [1, 2] and has been postulated to be associated with the capability to support universal computation [3, 4, 5].

In all of these relevant fields, distributed computation is generally discussed in terms of “memory”, “communication”, and “processing”. Memory refers to the storage of information by an agent or process to be used in its future. It has been investigated in coordinated motion in modular robots [6], in the dynamics of inter-event distribution times [7], and in synchronization between coupled systems [8]. Communication refers to the transfer of information between one agent or process and another; it has been shown to be of relevance to biological systems (e.g. dipole-dipole interaction in microtubules [9], and in signal transduction by calcium ions [10]), social animals (e.g. schooling behavior in fish [11]), and agent-based systems (e.g. the influence of agents over their environments [12], and in inducing emergent neural structure [13]). Processing refers to the combination of stored and/or transmitted information into a new form; it has been discussed in particular for biological neural

networks and models thereof [14, 15, 16, 17] (where it has been suggested as a potential biological driver), and also regarding collision-based computing (e.g. [18, 19], and including soliton dynamics and collisions [20]).

Significantly, these terms correspond to the component operations of Turing universal computation: *information storage*, *information transfer* (or transmission) and *information modification*. Yet despite the obvious importance of these *information dynamics*, we have no framework for either quantifying them individually or understanding how they interact to give rise to distributed computation. Here, we present the first complete framework which quantifies each of the information dynamics or component operations of computation on a *local scale* and describes how they inter-relate to produce distributed computation. Our focus on the local scale within the system is an important one. Several authors have suggested that a complex system is better characterized by studies of its local dynamics than by averaged or overall measures (e.g. [21, 22]), and indeed here we believe that quantifying and understanding distributed computation will necessitate studying the information dynamics and their interplay on a local scale in space and time. Additionally, we suggest that the quantification of the individual information dynamics of computation provides three *axes of complexity* within which to investigate and classify complex systems, allowing deeper insights into the variety of computation taking place in different systems.

An important focus for discussions on the nature of distributed computation have been cellular automata (CAs)

*jlizier@it.usyd.edu.au

as model systems offering a range of dynamical behavior, including supporting complex computations and the ability to model complex systems in nature [1]. We select CAs for experimentation here because there is very clear *qualitative* observation of emergent structures representing information storage, transfer and modification therein (e.g. [1, 3]). CAs are a critical proving ground for any theory on the nature of distributed computation: significantly, Von Neumann was known to be a strong believer that “a general theory of computation in ‘complex networks of automata’ such as cellular automata would be essential both for understanding complex systems in nature and for designing artificial complex systems” ([1] describing [23]).

Information theory provides the logical platform for our investigation, and we begin with a summary of the main information-theoretic concepts required. We provide additional background on the qualitative nature of distributed computation in CAs, highlighting the opportunity for our framework to provide quantitative insights here. Subsequently, we consider each component operation of universal computation in turn, and describe how to quantify it locally in a spatiotemporal system. As an application, we measure each of these information dynamics at every point in space-time in several important CAs. Our framework provides the first complete quantitative evidence for a well-known set of conjectures on the emergent structures dominating distributed computation in CAs: that blinkers provide information storage, particles provide information transfer, and particle collisions facilitate information modification. Furthermore, our results imply that the *coherence* of information may be a defining feature of complex distributed computation. Our findings are significant because these emergent structures of computation in CAs have known analogues in many physical systems (e.g. solitons and biological pattern formation processes), and as such this work will contribute to our fundamental understanding of the nature of distributed computation and the dynamics of complex systems.

II. INFORMATION-THEORETICAL PRELIMINARIES

Information theory is an obvious tool for quantifying the information dynamics involved in distributed computation. In fact, information theory has already proven to be a useful framework for the design and analysis of complex self-organized systems [24]. Here, we will extend this success to describing distributed computation in complex systems.

We begin by reviewing several necessary information theoretic quantities (generally following the formulation in [25]). The fundamental quantity is the Shannon *entropy*, which represents the uncertainty associated with any measurement x of a random variable X (units in

bits):

$$H_X = - \sum_x p(x) \log_2 p(x). \quad (1)$$

The *joint entropy* of two (or more) random variables X and Y is a generalization to quantify the uncertainty of the joint distribution of X and Y :

$$H_{X,Y} = - \sum_{x,y} p(x,y) \log_2 p(x,y). \quad (2)$$

The *conditional entropy* of X given Y is the average uncertainty that remains about x when y is known:

$$H_{X|Y} = - \sum_{x,y} p(x,y) \log_2 p(x|y). \quad (3)$$

The *mutual information* between X and Y measures the average reduction in uncertainty about x that results from learning the value of y , or vice versa:

$$I_{X;Y} = \sum_{x,y} p(x,y) \log_2 \frac{p(x,y)}{p(x)p(y)}. \quad (4)$$

$$I_{X;Y} = H_X - H_{X|Y} = H_Y - H_{Y|X}. \quad (5)$$

The *conditional mutual information* between X and Y given Z is the mutual information between X and Y when Z is known:

$$I_{X;Y|Z} = H_{X|Z} - H_{X|Y,Z} \quad (6)$$

$$= H_{Y|Z} - H_{Y|X,Z}. \quad (7)$$

The *entropy rate* is the limiting value of the rate of change of the joint entropy over k consecutive states of X , (i.e. measurements $x^{(k)}$ of the random variable $X^{(k)}$), as k increases [26]:

$$H_{\mu X} = \lim_{k \rightarrow \infty} \frac{H_{X^{(k)}}}{k} = \lim_{k \rightarrow \infty} H'_{\mu X}(k), \quad (8)$$

$$H'_{\mu X}(k) = \frac{H_{X^{(k)}}}{k}. \quad (9)$$

The entropy rate can also be expressed as the limiting value of the conditional entropy of the next state of X (i.e. measurements x_{n+1} of the random variable X') given knowledge of the previous k states of X (i.e. measurements $x_n^{(k)}$, up to and including time step n , of the random variable $X^{(k)}$):

$$H_{\mu X} = \lim_{k \rightarrow \infty} H_{X'|X^{(k)}} = \lim_{k \rightarrow \infty} H_{\mu X}(k), \quad (10)$$

$$H_{\mu X}(k) = H_{X^{(k+1)}} - H_{X^{(k)}}. \quad (11)$$

Grassberger [27] first noticed that a slow approach of the entropy rate to its limiting value was a sign of complexity. Formally, Crutchfield and Feldman [26] use the conditional entropy form of the entropy rate (10)[69] to observe that at a finite block size k , the difference

$H_{\mu X}(k) - H_{\mu X}$ represents the information carrying capacity in size k -blocks that is due to correlations. The sum over all k gives the total amount of structure in the system, quantified as excess entropy (measured in bits):

$$E_X = \sum_{k=0}^{\infty} [H_{\mu X}(k) - H_{\mu X}]. \quad (12)$$

The excess entropy can also be formulated as the mutual information between the semi-infinite past and semi-infinite future of the system:

$$E_X = \lim_{k \rightarrow \infty} I_{X^{(k)}; X^{(k+)}}, \quad (13)$$

where $X^{(k+)}$ is the random variable (with measurements $x_{n+1}^{(k+)}$) referring to the k future states of X (from time step $n + 1$ onwards). This interpretation is known as the *predictive information* [28], as it highlights that the excess entropy captures the information in a system's past which is relevant to predicting its future.

III. CELLULAR AUTOMATA

A. Introduction to Cellular Automata

Cellular automata (CA) are discrete dynamical systems consisting of an array of cells which each synchronously update their discrete state as a function of the states of a fixed number of spatially neighboring cells using a uniform rule. Although the behavior of each individual cell is very simple, the (non-linear) interactions between all cells can lead to very intricate global behavior meaning CAs have become a classic example of self-organized complex behavior. Of particular importance, CAs have been used to model real-world spatial dynamical processes, including fluid flow, earthquakes and biological pattern formation [1].

The neighborhood of a cell used as inputs to its update rule at each time step is usually some regular configuration. In 1D CAs, this means the same range r of cells on each side and including the current state of the updating cell. One of the simplest variety of CAs – 1D CAs using binary states, deterministic rules and one neighbor on either side – are known as the *Elementary CAs*, or *ECAs*. Example evolutions of ECAs from random initial conditions may be seen in Fig. 2(a) and Fig. 6(a). For more complete definitions of CAs, including the definition of the Wolfram rule number convention for specifying update rules, see [29].

Wolfram [4, 29] sought to classify the asymptotic behavior of CA rules into four classes: I. Homogeneous state; II. Simple stable or periodic structures; III. Chaotic aperiodic behavior; IV. Complicated localized structures, some propagating. Much conjecture remains as to whether these classes are quantitatively distinguishable (e.g. see [30]), however they do provide an interesting analogy (for discrete state and time) to our knowledge

of dynamical systems, with classes I and II representing ordered behavior, class III representing chaotic behavior, and class IV representing complex behavior and considered as lying between the ordered and chaotic classes.

More importantly, the approach seeks to characterize complex behavior in terms of emergent structure in CAs, surrounding *gliders*, *particles* and *domains*. Qualitatively, a domain may be described as a set of background configurations in a CA, for which any given configuration will update to another such configuration in the set in the absence of any disturbance. Domains are formally defined within the framework of computational mechanics [22] as spatial process languages in the CA. Particles are qualitatively considered to be moving elements of coherent spatiotemporal structure. Gliders are particles which repeat periodically in time while moving spatially (repetitive non-moving structures are known as *blinkers*). Formally, particles are defined within the framework of computational mechanics as a boundary between two domains [22]; as such, they can also be termed as domain walls, though this is typically used with reference to aperiodic particles.

These emergent structures are more clearly visible when the CA is filtered in some way, using for example ϵ -machines [22], input entropy [31], local information [32], or local statistical complexity [21]. All of these filtering techniques produce a *single* filtered view of the structures in the CA: our measures of local information dynamics will present several filtered views of the distributed computation in a CA. The ECA examples analyzed in this paper are introduced in Section III C.

B. Computation in Cellular Automata

CAs can be interpreted as undertaking distributed computation: it seems fairly clear that “data represented by initial configurations is processed by time evolution” [4]. As such, computation in CAs has been a popular topic for study (see [1]), with a particular focus in observing or constructing (Turing) universal computation in certain CAs. An ability for universal computation is defined to be where “suitable initial configurations can specify arbitrary algorithm procedures” in the computing entity, which is capable of “evaluating any (computable) function” [4]. Wolfram conjectured that all class IV complex CAs were capable of universal computation [4, 33]. He went on to state that prediction in systems exhibiting universal computation is limited to explicit simulation of the system, as opposed to the availability of any simple formula or “short-cut”, drawing parallels to the halting problem for universal Turing machines [4, 33] which are echoed by Langton [3] and Casti [5]. (Casti extended the analogy to undecidable statements in formal systems, i.e. Gödel’s Theorem). The capability for universal computation has been proven for several CA rules, through the design of rules generating elements to (or by identifying elements which) specifically provide the component op-

erations required for universal computation: information storage, transmission and modification. Examples here include most notably the Game of Life [34] and ECA rule 110 [35]; also see [36] and discussions in [1].

The focus on elements providing information storage, transmission and modification pervades discussion of all types of computation in CAs (e.g. also see [19, 37]). Wolfram claimed that in class III CAs information propagates over an infinite distance at a finite speed, while in class IV CAs information propagates irregularly over an infinite range [33]. Langton [3] hypothesized that complex behavior in CAs exhibited the three component operations required for universal computation. He suggested that the more chaotic a system becomes the more information transmission increases, and the more ordered a system becomes the more information it stores. Complex behavior was said to occur at a phase transition between these extremes requiring an intermediate level of both information storage and transmission: if information propagates too well, *coherent information decays into noise*. Langton elaborates that transmission of information means that the “dynamics must provide for the propagation of information in the form of signals over arbitrarily long distances”, and suggests that particles in CAs form the basis of these signals. To complete the qualitative identification of the elements of computation in CAs, he also suggested that blinkers formed the basis of information storage, and collisions between propagating (particles) and static structures (blinkers) “can modify either stored or transmitted information in the support of an overall computation”. Rudimentary attempts were made at quantifying the average information transfer (and to some extent information storage), via mutual information (although as discussed later this is a symmetric measure not capturing directional transfer). Recognizing the importance of the emergent structures to computation, several examples exist of attempts to automatically identify CA rules which give rise to particles and gliders, e.g. [31, 38], suggesting these to be the most interesting and complex CA rules.

Several authors however criticize the aforementioned approaches of attempting to classify CAs in terms of their generic behavior or “bulk statistical properties”, suggesting that the wide range of differing dynamics taking place across the CA makes this problematic [1, 22]. Gray suggests that there may indeed be classes of CAs capable of more complex computation than universal computation alone [30]. More importantly, Hanson and Crutchfield [22] criticize the focus on universal computational ability as drawing away from the ability to identify “generic computational properties”, i.e. a lack of ability for universal computation does not mean a CA is not undertaking any computation at all. Alternatively, these studies suggest that analyzing the rich space-time dynamics *within* the CA is a more appropriate focus. As such, these and other studies have analyzed the *local* dynamics of intrinsic or other specific computation, focusing on particles facilitating the transfer of informa-

tion and collisions facilitating the information processing. Noteworthy examples here include: the method of applying filters from the domain of computational mechanics by Hanson and Crutchfield [22]; and analysis using such computational mechanics filters of CA rules selected via evolutionary computation to perform classification tasks by Mitchell et al [39, 40]. Related are studies which deeply investigate the nature of particles and their interactions (e.g. particle types and their interaction products identified for particular CAs in [40, 41, 42], and rules established for their interaction products in [43]).

Despite such interest, there is no complete framework that locally quantifies the individual information dynamics of distributed computation within CAs or other systems. In this study, we outline how the information dynamics can be locally quantified within the spatiotemporal structure of a CA. In particular, we describe the dynamics of how information storage and information transfer interact to give rise to information processing. Our approach is not to quantify computation or overall complexity, nor to identify universal computation or determine what is being computed; it is simply intended to quantify the component operations in space-time.

C. Examples of distributed computation in CAs

In this paper, we will examine the computation carried out by several important ECA rules:

- Class IV complex rules 110 and 54 [29] (see Fig. 4(a) and Fig. 2(a)), both of which exhibit a number of glider types and collisions. ECA rule 110 is the only proven computationally universal ECA rule [35].
- Rules 22 and 30 as representative class III chaotic rules [29] (see rule 22 in Fig. 7(a));
- Rules 18 as a class III rule which contains domain walls against a chaotic background domain [22, 44].

These CAs each carry out an *intrinsic* computation of the evolution to their ultimate attractor and phase on it (see [31] for a discussion of attractors and state space in finite-sized CAs).

We also examine a CA carrying out a “human-understandable” computational task. ϕ_{par} is a 1D CA with range $r = 3$ (Wolfram rule number 0xfeedffdec1aaec0eef000a0e1a020a0) that was evolved by Mitchell et al [39, 40] to classify whether the initial CA configuration had a majority of 1’s or 0’s by reaching a fixed-point configuration of all 1’s for the former or all 0’s for the latter. This CA rule achieved a success rate above 70% in its task. An example evolution of this CA can be seen in Fig. 5(a). The CA appears to carry out this computation using blinkers and domains for information storage, gliders for information transfer and glider collisions for information modification. The CA exhibits an initial emergence of domain regions of all 1’s or all 0’s storing information about local high densities of either

state. Where these domains meet, a checkerboard domain propagates slowly (1 cell per time step) in both directions, transferring information of a *soft* uncertainty in this part of the CA. Some “certainty” is provided where the checkerboard encounters a blinker boundary between 0 and 1 domains, which stores information about a *hard* uncertainty in that region of the CA. This results in an information modification event where the domain on the opposite side of the blinker to the incoming checkerboard is concluded to represent the higher density state, and is allowed to propagate over the checkerboard. Because of the greater certainty attached to this decision, this new information transfer occurs at a faster speed (3 cells per time step); it can overrun checkerboard regions, and in fact collisions of opposing types of this strong propagation give rise to the (hard uncertainty) blinker boundaries in the first place. The final configuration is therefore the result of this distributed computation.

Quantification of the local information dynamics via these *three axes of complexity* (information storage, transfer and modification) will provide quite detailed insights into the distributed computation carried out in a system. In all of these CAs we expect local measures of information storage to highlight blinkers and domain regions, local measures of information transfer to highlight particles (including gliders and domain walls), and local measures of information modification to highlight particle collisions.

This will provide a deeper understanding of computation than single or generic measures of bulk statistical behavior, from which conflict often arises in attempts to provide classification of complex behavior. In particular, we seek clarification on the long-standing debate regarding the nature of computation in ECA rule 22. Suggestions that rule 22 is complex include the difficulty in estimating the metric entropy (i.e. temporal entropy rate) for rule 22 in [27], due to “complex long-range effects, similar to a critical phenomenon” [45]. This effectively corresponds to an implication that rule 22 has contains an infinite amount of memory (see Section IV A). Also, from an initial condition of only a single “on” cell, rule 22 forms a pattern known as the “Sierpinski Gasket” [29] which exhibits clear self-similar structure. Furthermore, rule 22 is a 1D mapping of the 2D Game of Life CA (known to have the capability for universal computation [34]) and in this sense is referred to as “life in one dimension” [46], and complex structure in the language generated by iterations of rule 22 has been identified [47]. Also, we report here that we have investigated the C_1 complexity measure [48] (an enhanced version of the variance of the input entropy [31]) for all ECAs, and found rule 22 to clearly exhibit the largest value of this metric (0.78 bits to rule 110’s 0.085 bits). On the other hand, suggestions that rule 22 is not complex include its high sensitivity to initial conditions leading to Wolfram classifying it as class III chaotic [29]. Gutowitz and Domain [49] claim this renders it as chaotic despite the subtle long-range effects it displays, further identifying its fast statistical

convergence, and exponentially long and thin transients in state space (see [31]). Importantly, no coherent structure (particles, collisions, etc.) is found for rule 22 using a number of known filters for such structure (e.g. local statistical complexity [21]): this reflects the paradigm shift to an examination of local dynamics rather than generic, overall or averaged analysis. In our approach, we seek to combine this local viewpoint of the dynamics with a quantitative breakdown of the individual elements of computation, and will investigate rule 22 in this light.

IV. INFORMATION STORAGE

In this section we outline methods to quantify information storage. We describe how total information storage is captured by excess entropy, and introduce active information storage to capture the amount of information storage that is currently in use. We present the first application of local profiles of both measures to cellular automata.

A. Excess entropy as total information storage

Although discussion of information storage or memory in CAs has often focused on periodic structures (particularly in construction of universal Turing machines), information storage does not necessarily entail periodicity. The excess entropy (Eq. (12,13)) more broadly encompasses all types of structure and memory by capturing correlations across all lengths of time, including non-linear correlations. It is quite clear from the predictive information formulation of the excess entropy Eq. (13) – as the information from a system’s past that is relevant to predicting its future – that it is a measure of the total information storage in a system.

We use the term *single-agent excess entropy* to refer to measuring the excess entropy for individual agents or cells using their one-dimensional time series of states. This is a measure of the *average* memory for *each agent*. Furthermore, we use the term *collective excess entropy* to refer to measuring the temporal excess entropy for a collective of agents (e.g. a set of neighboring cells in a CA) using their two-dimensional time series of states. Considered as the mutual information between their joint past and future (i.e. a joint temporal predictive information), this is a measure of the *average* total memory stored in the collective (i.e. stored *collectively* by a set of cells in a CA). Collective excess entropy could be used for example to quantify the “undiscovered *collective memory* that may present in certain fish schools” [11].

Grassberger studied temporal entropy rate estimates for several ECAs in [27, 45] in order to gain insights into their excess entropies. These studies estimated temporal entropy rates $H_{\mu,N}(k)$ for spatial blocks of size N as N is increased. Estimated values of $H_{\mu,N}(k)$ (for $N = 1$ and in the limit as $N \rightarrow \infty$) were cataloged for most ECAs

in a table of statistical properties in [50]. Using these estimates, the studies focused on inferring the collective excess entropies rather than the single-agent excess entropies ($N = 1$). For several rules (including rule 22, studied with Monte Carlo estimates), the temporal entropy rate estimates $H_{\mu,N}(k)$ (for $N > 1$) were concluded to follow a power law decay to their asymptote $H_{\mu,N}$:

$$H_{\mu,N}(k) = H_{\mu,N} + C/k^\alpha, \quad (14)$$

with exponent $\alpha \leq 1$ (C is a constant). This is significant because with $\alpha \leq 1$ the collective excess entropy (known as *effective measure complexity* in [27]) is divergent, implying a highly complex process. This case has been described as “a phenomenon which can occur in more complex environments”, as with strong long-range correlations a semi-infinite sequence “could store an infinite amount of information about its continuation” [51] (as per the predictive information form of the excess entropy Eq. (13)). Rule 22 was inferred to have $H_{\mu,N} = 0$ and infinite excess entropy, which can be interpreted as a process requiring an infinite amount of memory to maintain an aperiodicity [26]. Alternative methods for computing two-dimensional excess entropies, which would be applicable for computing the collective excess entropy in CAs, were presented by Feldman and Crutchfield in [52].

In attempting to quantify *local* information dynamics of distributed computation here, our focus is on information storage for *single agents or cells* rather than the joint information storage across the collective. Were such power-law trends to exist for the single-agent case, they may be more significant than for the collective case: divergent collective excess entropy implies that the collective is at least trivially utilizing all of its available memory (and even the chaotic rule 30 exhibits this), whereas divergent single-agent excess entropy implies that all agents are individually highly utilizing the resources of the collective in a highly complex process. One could go on to study the entropy rate convergence for single agents ($N = 1$) [70], however any findings would be subject to the problems with overall or averaged measures described earlier. Again, we emphasize that our focus is on local measures in *time* as well as *space*, which we present in the next section.

First though we note that with respect to CAs, where each cell has only a finite number of states b and takes direct influence from only its single past state and the states of a finite number of neighbors, the meaning of (either average or local) information storage being greater $\log_2 b$ bits (let alone infinite) in the state “process” of a single cell is not immediately obvious. Clearly, a cell in an ECA cannot store more than 1 bit of information in isolation. However, the *bidirectional* communication in CAs effectively allows a cell to store extra information in neighbors (even beyond the immediate neighbors), and to subsequently retrieve that information from those neighbors at a later point in time. While measurement of the excess entropy does not explicitly look for such *self-influence* communicated through neighbors, it is indeed

the method by which a significant portion of information is channeled. Considering the predictive information interpretation in Eq. (13), it is easy to picture self-influence between semi-infinite past and future blocks being conveyed via neighbors (see Fig. 1(a)). This is akin to the use of stigmergy (indirect communication through the environment, e.g. see [53]) to communicate with oneself.

A measurement of more than $\log_2 b$ bits stored by a cell on average, or indeed an infinite information storage, is then a perfectly valid result: in an infinite CA, each cell has access to an infinite amount of neighbors in which to store information which can later be used to influence its own future. Note however, that since the storage medium is shared by all cells, one should not think about the total memory as the total number of cells multiplied by this average. The total memory would be properly measured by the collective excess entropy, which takes into account the inherent redundancy here.

B. Local excess entropy

As discussed previously, we now shift focus to local measures of information storage, which have the potential to provide more detailed insights into information storage structures and their involvement in computation than single ensemble measures.

The *local excess entropy* is a measure of how much information a given agent is *currently* storing at a particular point in time. To derive it, note that the excess entropy of a process is actually the *expectation value* of the local excess entropy for the process at every time step [54]. [71] The local excess entropy $e_X(n+1)$ of a process is simply the log term from inside the mutual information expansion (as per Eq. (4)) of the predictive information formulation of excess entropy in Eq. (13), evaluated for the semi-infinite past and future at the given time step $n+1$:

$$e_X(n+1) = \lim_{k \rightarrow \infty} \log_2 \frac{p(x_n^{(k)}, x_{n+1}^{(k+)})}{p(x_n^{(k)})p(x_{n+1}^{(k+)})}. \quad (15)$$

Note that the excess entropy is the average of the local values, $E_X = \langle e_X(n) \rangle$, and that by convention we use lower-case symbols to denote local values. The limit $k \rightarrow \infty$ is an important part of this definition, since correlations at all time scales should be included in the computation of information storage. Since this is not computationally feasible in general, we retain the notation $e_X(n+1, k)$ to denote finite- k estimates of $e_X(n+1)$.

The notation is generalized for lattice systems (such as CAs) with *spatially-ordered* agents to represent the local excess entropy for cell X_i at time $n+1$ as:

$$e(i, n+1) = \lim_{k \rightarrow \infty} \log_2 \frac{p(x_{i,n}^{(k)}, x_{i,n+1}^{(k+)})}{p(x_{i,n}^{(k)})p(x_{i,n+1}^{(k+)})}. \quad (16)$$

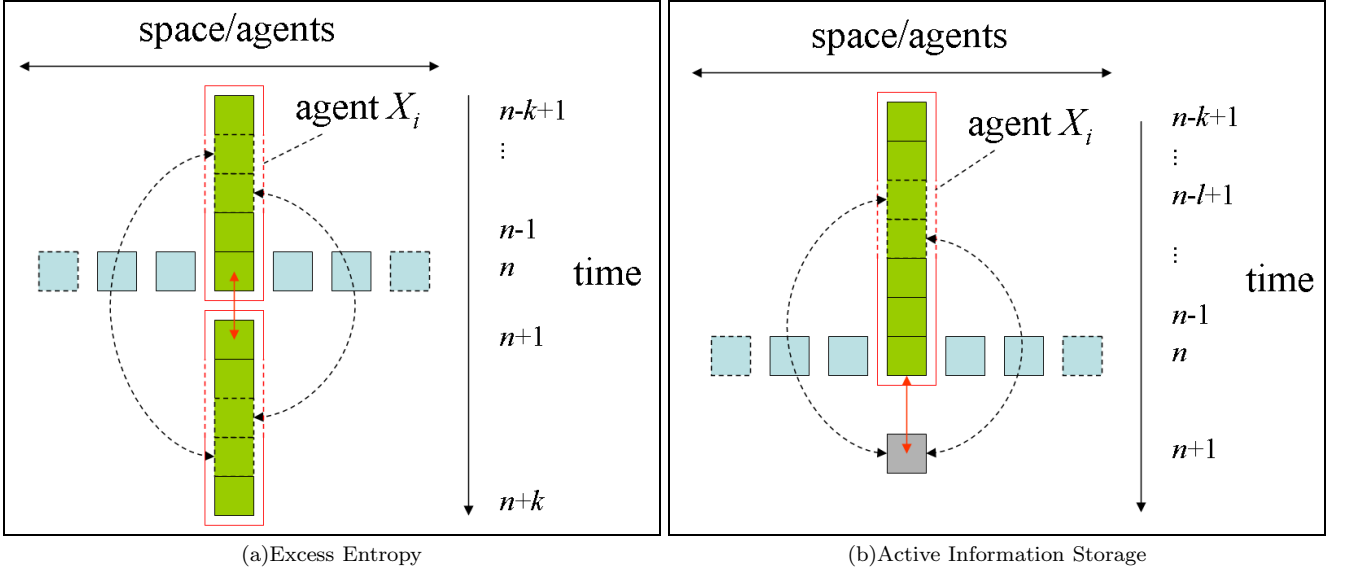


FIG. 1: Measures of single-agent information storage in distributed systems. (a) Excess Entropy: *total* information from the cell's past that is relevant to predicting its future. (b) Active Information storage: the information storage that is *currently in use* in determining the next state of the cell. The stored information can be conveyed directly through the cell itself or via neighboring cells.

Again, $e(i, n+1, k)$ is used to denote finite- k estimates of $e(i, n+1)$. Local excess entropy is defined for every spatiotemporal point (i, n) in the system. (Alternatively, the collective excess entropy can only be localized in time).

While the average excess entropy is always positive, the local excess entropy may in fact be positive or negative, meaning the past history of the cell can either positively inform us or actually misinform us about its future. An observer is misinformed where the semi-infinite past and future are relatively unlikely to be observed together as compared to their independent likelihoods.

C. Active information storage

The excess entropy measures the total stored information which will be used *at some point* in the future of the state process of an agent, possibly but not necessarily at the next time step $n+1$. In examining the local information dynamics of computation, we are interested in how much of the stored information is actually *in use* at the next time step. As we will see in Section VI, this is particularly important in understanding how stored information interacts with information transfer in information processing. As such, we derive *active information storage* A_X as the average mutual information between the semi-infinite past of the process and its next state, as opposed to its whole (semi-infinite) future:

$$A_X = \lim_{k \rightarrow \infty} I(X^{(k)}; X'). \quad (17)$$

The *local* active information storage is then a measure of the amount of information storage in use by the process

at a particular time-step $n+1$:

$$a_X(n+1) = \lim_{k \rightarrow \infty} \log_2 \frac{p(x_n^{(k)}, x_{n+1})}{p(x_n^{(k)})p(x_{n+1})}, \quad (18)$$

and we have $A_X = \langle a_X(n) \rangle$. We retain the notation $a_X(n+1, k)$ for finite- k estimates. Again, we generalize the measure for agent X_i in a lattice system as:

$$a(i, n+1) = \lim_{k \rightarrow \infty} \log_2 \frac{p(x_{i,n}^{(k)}, x_{i,n+1})}{p(x_{i,n}^{(k)})p(x_{i,n+1})}, \quad (19)$$

and use $a(i, n+1, k)$ to denote finite- k estimates there, noting that the local active information storage is defined for every spatiotemporal point (i, n) in the lattice system.

The average active information storage will always be positive (as for the excess entropy), but is bounded above by $\log_2 b$ bits where the agent only takes b discrete states. The local active information storage is not bound in this manner however, with values larger than $\log_2 b$ indicating that the particular past of an agent provides strong positive information about its next state. Furthermore, the local active information storage can be negative, where the past history of the agent is actually misinformative about its next state. An observer is misinformed where the past history and observed next state are relatively unlikely to occur together as compared to their separate occurrence.

D. Local information storage results

In order to evaluate the local measures within sample CA runs, we estimate the required probability distribu-

tion functions from CA runs of 10 000 cells, initialized from random states, with 600 time steps retained (after the first 30 time steps were eliminated to allow the CA to settle). Alternatively, for ϕ_{par} we used 1000 cells with 1000 time steps retained. Periodic boundary conditions were used. Observations taken at every spatiotemporal point in the CA were used in estimating the required probability distribution functions, since the cells in the CA are homogeneous agents. All results were confirmed by at least 10 runs from different initial states.

We make estimates of the measures with finite values of k , noting that the insights described here could not be attained unless a reasonably large value of k was used in order to capture a large proportion of the correlations. Determination of an appropriate value of k was discussed in [55] for the related transfer entropy measure, presented in Section V. As a rule of thumb, k should at least be larger than the period of any regular background domain in order to capture the information storage underpinning its continuation.

We begin by examining the results for rules 54 and 110, which contain regular gliders against periodic background domains. For the CA evolutions in Fig. 2(a) and Fig. 4(a), the local profiles of $e(i, n, k = 8)$ generated are displayed in Fig. 2(b) and Fig. 4(b) (positive values only), and the local profiles of $a(i, n, k = 16)$ in Fig. 2(c) and Fig. 4(c) (positive values only). It is quite clear that positive information storage is concentrated in the vertical gliders or blinkers, and the domain regions. As expected, these results provide quantitative evidence that the blinkers are the dominant information storage entities. That the domain regions contain significant information storage should not be surprising, since as a periodic sequence its past does indeed store information about its future. In fact, the local values for each measure form spatially and temporally periodic patterns in the domains, due to the spatial and temporal periodicities exhibited there. While the local active information storage indicates a similar amount of stored information in use to compute each space-time point in both the domain and blinker areas, the local excess entropy reveals a larger *total* amount of information is stored in the blinkers. For the blinkers known as α and β in rule 54 [43] this is because the temporal sequences of the center columns of the blinkers (0-0-0-1, with $e(i, n, k = 8)$ in the range 5.01 to 5.32 bits) are more complex than those in the domain (0-0-1-1 and 0-1, with $e(i, n, k = 8)$ in the range 1.94 to 3.22 bits), even where they are of the same period. We have $e(i, n, k = 8) > 1$ bit here due to the distributed information storage supported by bidirectional communication (as discussed earlier): this also supports the period-7 domain in rule 110. Another area of strong information storage appears to be the “wake” of the more complex gliders in rule 110 (see the glider at top left of Fig. 4(b) and Fig. 4(c)). This result aligns well with our observation in [55] that the dynamics following the leading edge of regular gliders consists largely of “non-traveling” information. The presence of the information

storage is shown by both measures, although the relative strength of the total information storage is again revealed only by the local excess entropy.

Negative values of $a(i, n, k = 16)$ for rules 54 and 110 are displayed in Fig. 2(d) and Fig. 4(d). Interestingly, negative local components of local active information storage measure are concentrated in the traveling glider areas (e.g. γ^+ and γ^- for rule 54 [43]), providing a good spatiotemporal filter the glider structure. This is because when a traveling glider is encountered at a given cell, the past history of that cell (being part of the background domain) is misinformative about the next state, since the domain sequence was more likely to continue than be interrupted. For example, see the marked positions of the γ gliders in Fig. 3. There we have $p(x_{n+1}|x_n^{(k=16)}) = 0.25$ and $p(x_{n+1}) = 0.52$: since the next state occurs relatively infrequently after the given history, we have a misinformative $a(n, k = 16) = -1.09$ bits. This is juxtaposed with the points four time steps before those marked “x”, which have the same history $x_n^{(k=16)}$ but are part of the domain, with $p(x_{n+1}|x_n^{(k=16)}) = 0.75$ and $p(x_{n+1}) = 0.48$ giving $a(n, k = 16) = 0.66$ bits, quantifying the positive information storage there. Note that the points with misinformative information storage are not necessarily those selected by other filtering techniques as part of the gliders: e.g. the finite state transducers technique (using left to right scanning by convention) [56] would identify points 3 cells to the right of those marked “x” as part of the γ^+ glider.

The local excess entropy produced some negative values around traveling gliders (results not shown), though these were far less localized on the gliders themselves and less consistent in occurrence than for the local active information storage. This is because the local excess entropy, as measure of total information storage into the future, is more loosely tied to the dynamics at the given spatiotemporal point. The effect of a glider encounter on $e(i, n, k)$ is smeared out in time, and in fact the dynamics may store more positive information in total than the misinformation encountered at the specific location of the glider. (For example, glider pairs in Fig. 4(b) have positive total information storage, since a glider encounter becomes much more likely in the wake of a previous glider).

As another rule containing regular gliders against a periodic background domain, analysis of the raw states of ϕ_{par} in Fig. 5(a) provides similar results for $e(i, n, k = 5)$ in Fig. 5(b) and $a(i, n, k = 10)$ in Fig. 5(c) and Fig. 5(d) here. One distinction is that the blinker here contains no more stored information than the domain, since it is no more complicated. Importantly, we confirm the information storage capability of the blinkers and domains in this human understandable computation.

Another interesting example is provided by ECA rule 18, which contains domain walls against a seemingly irregular background domain. We measured the local values for $e(i, n, k = 8)$ and $a(i, n, k = 16)$ (see Fig. 6(b), Fig. 6(d) and Fig. 6(e)) for the raw states of rule 18 dis-

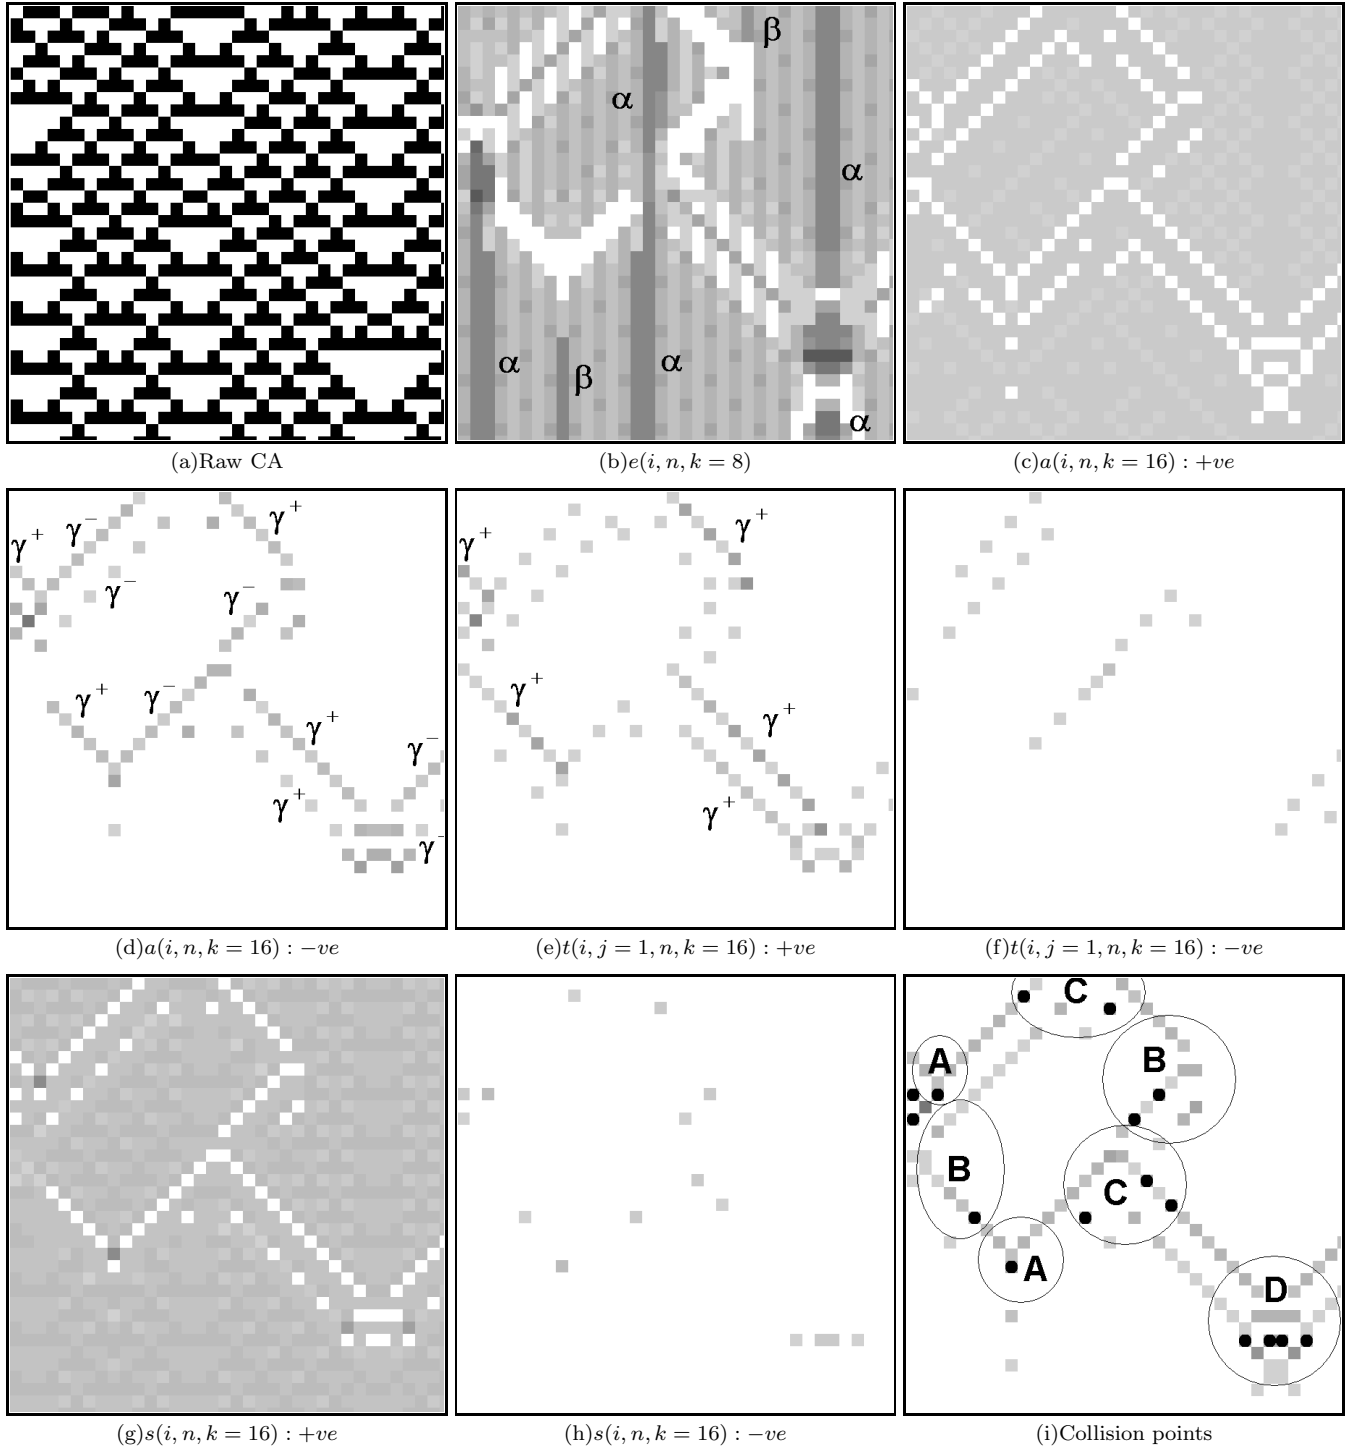


FIG. 2: Local information dynamics in rule 54: (35 time steps displayed for 35 cells, time increases down the page for all CA plots): (b) Local excess entropy, positive values only, (all figures gray-scale with 16 levels) with max. 11.79 bits (black), min. 0.00 bits (white); Local active information: (c) positive values only, max. 1.07 bits (black), min. 0.00 bits (white), (d) negative values only, max. 0.00 bits (white), min. -12.27 bits (black); Local apparent transfer entropy (one cell to the right): (e) positive values only, max. 7.93 bits (black), min. 0.00 bits (white), (f) negative values only, max. 0.00 bits (white), min. -4.04 bits (black); Local separable information: (g) positive values only, max. 8.40 bits (black), min. 0.00 bits (white), (h) negative values only, max. 0.00 bits (white), min. -5.27 bits (black); (i) Positions of $s(i, n, k = 16) < 0$ marked against the local collective transfer entropy profile $t(i, n, k = 16)$.

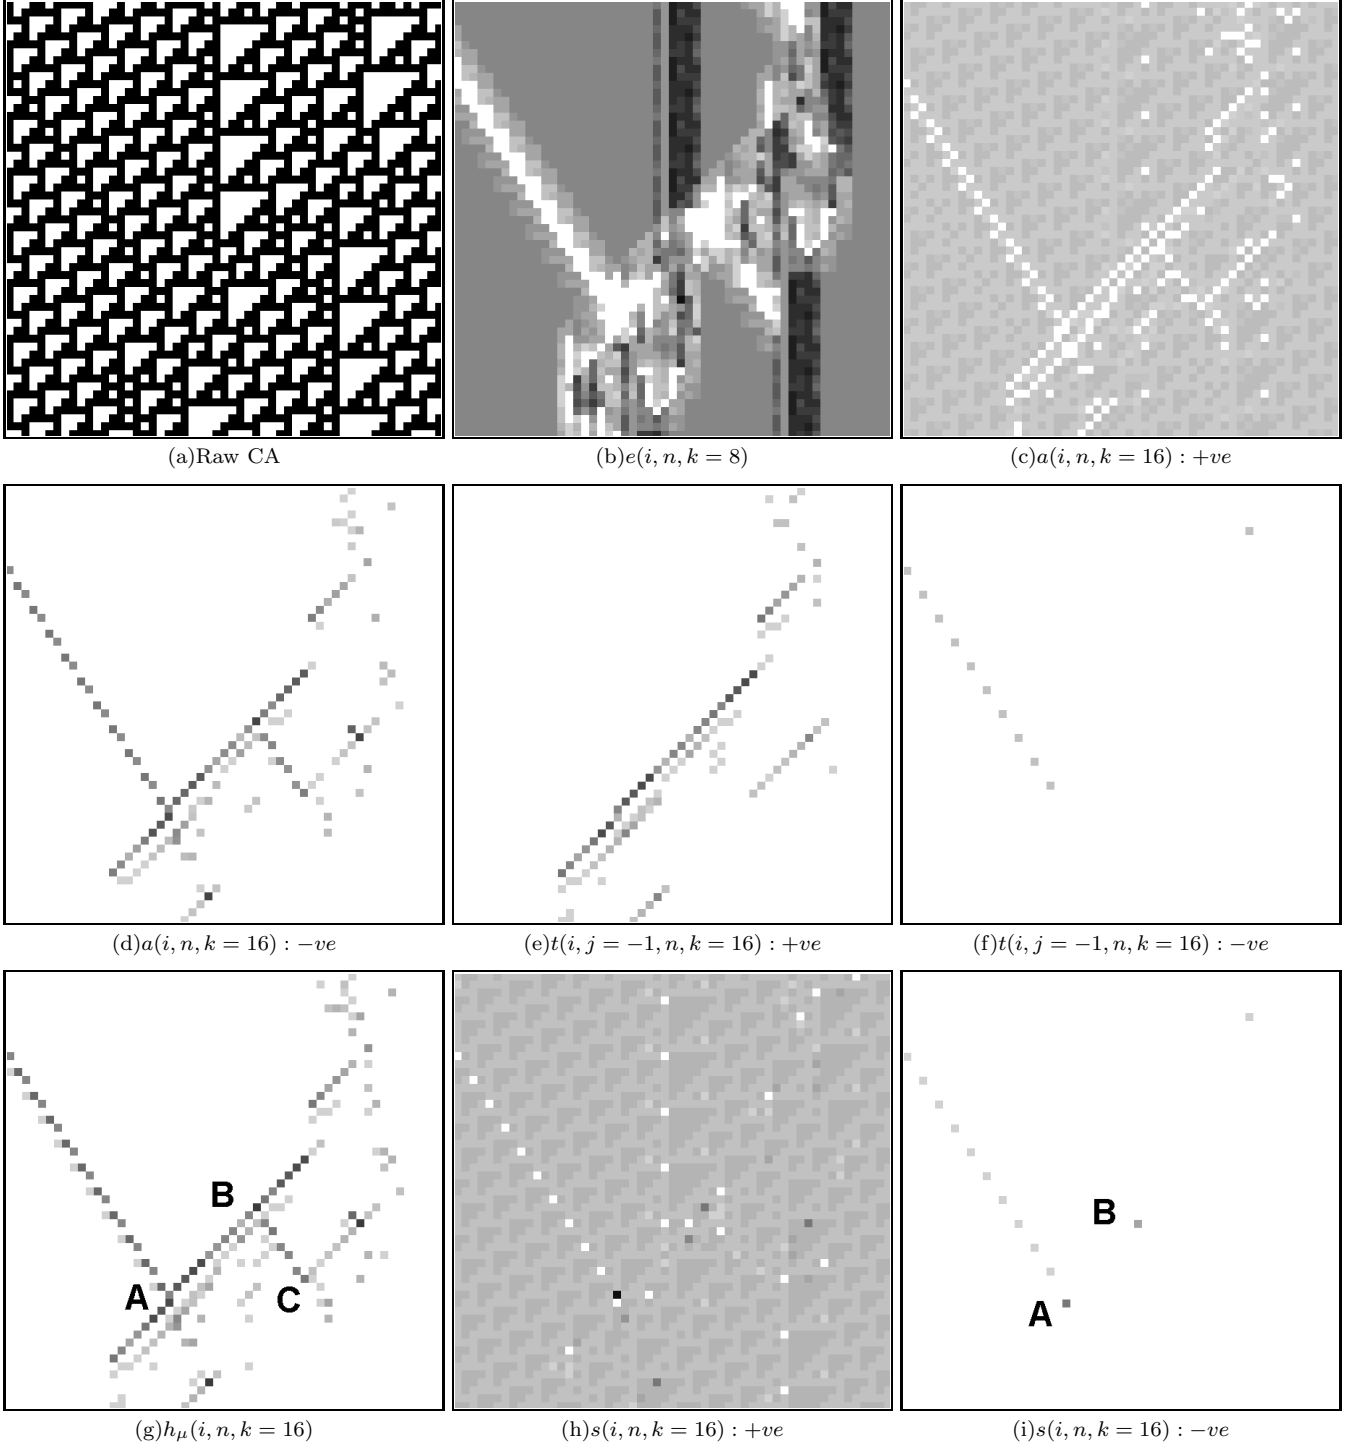


FIG. 4: Local information dynamics in rule 110: (55 time steps displayed for 55 cells): (b) Local excess entropy, positive values only, max. 10.01 bits (black), min. 0.00 bits (white); Local active information: (c) positive values only, max. 1.22 bits (black), min. 0.00 bits (white), (d) negative values only, max. 0.00 bits (white), min. -9.21 bits (black); Local apparent transfer entropy (one cell to the left): (e) positive values only, max. 10.43 bits (black), min. 0.00 bits (white), (f) negative values only, max. 0.00 bits (white), min. -6.01 bits (black); (g) Local temporal entropy rate, max. 10.43 bits (black), min. 0.00 bits (white); Local separable information: (h) positive values only, max. 5.47 bits (black), min. 0.00 bits (white), (i) negative values only, max. 0.00 bits (white), min. -5.20 bits (black).

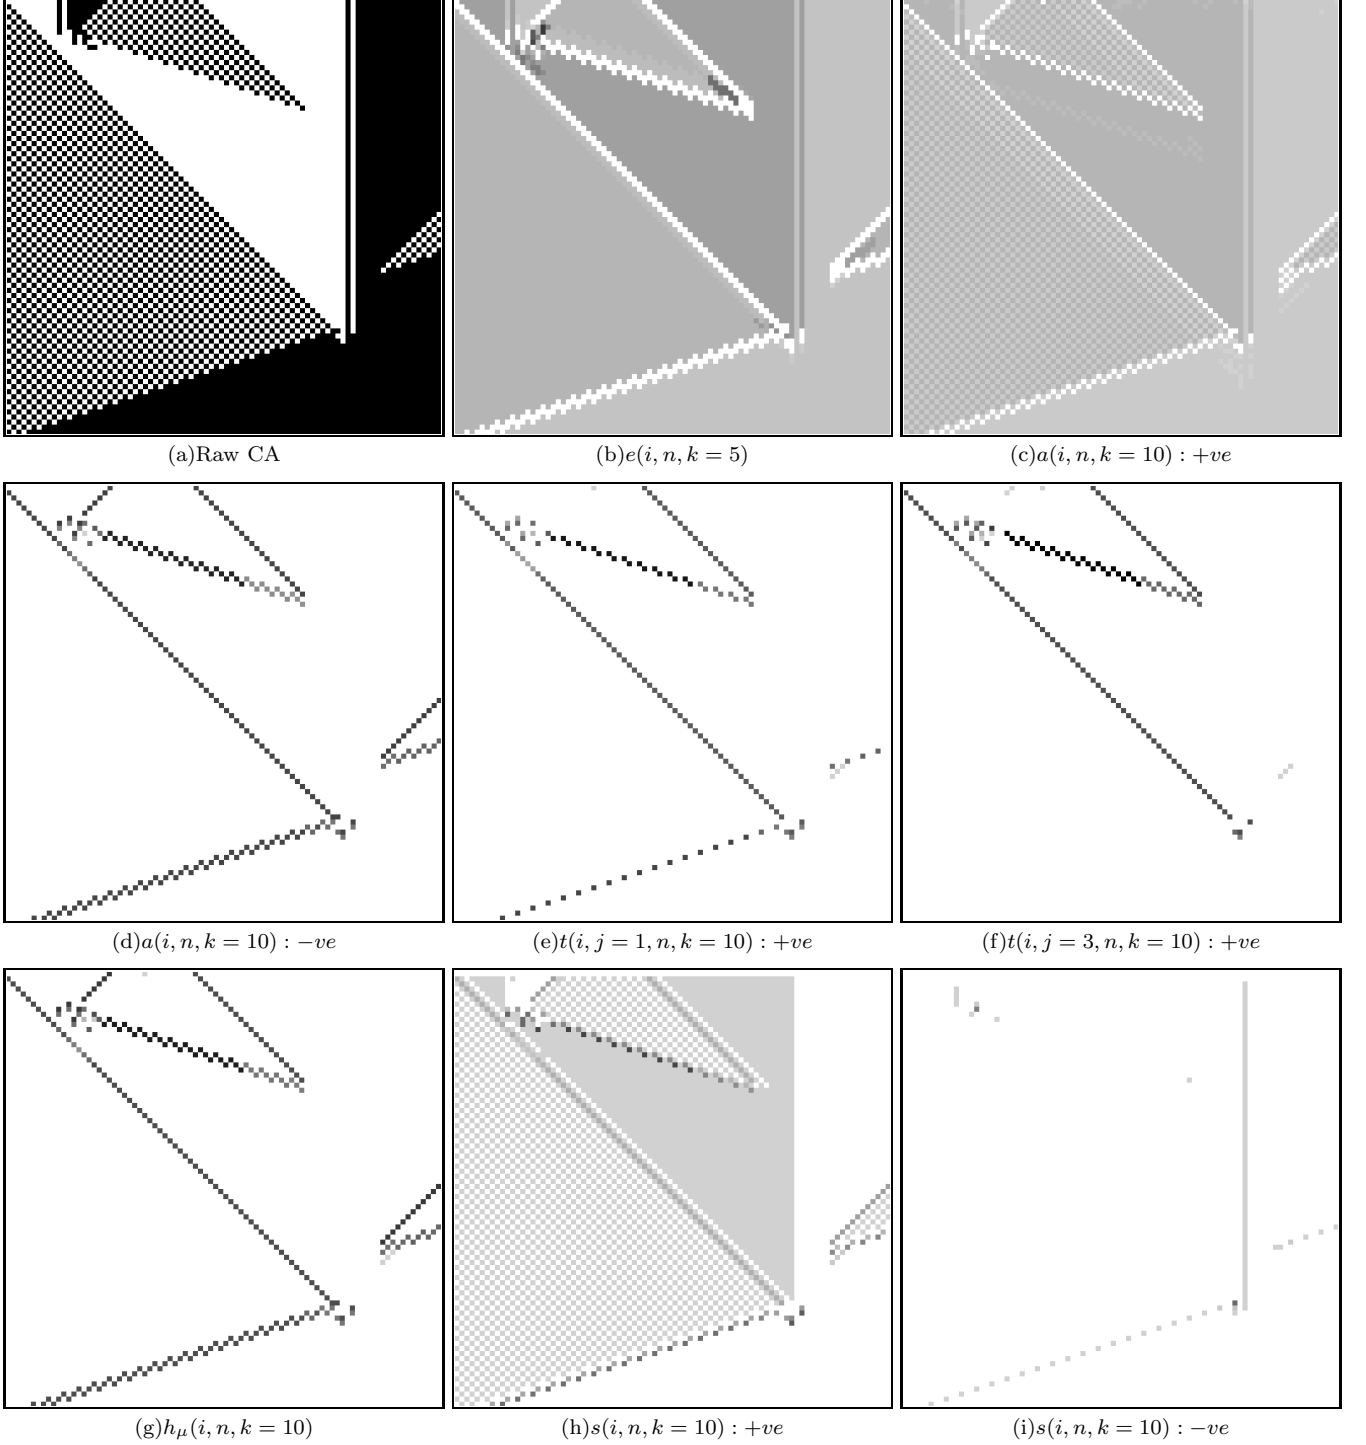


FIG. 5: Local information dynamics in $r = 3$ rule ϕ_{par} : (86 time steps displayed for 86 cells): (b) Local excess entropy, positive values only, max. 11.76 bits (black), min. 0.00 bits (white); Local active information: (c) positive values only, max. 1.52 bits (black), min. 0.00 bits (white), (d) negative values only, max. 0.00 bits (white), min. -9.41 bits (black); Local apparent transfer entropy: (e) one cell to the right, positive values only, max. 10.45 bits (black), min. 0.00 bits (white), (f) three cell to the right, positive values only, max. 9.24 bits (black), min. 0.00 bits (white); (g) Local temporal entropy rate, max. 10.92 bits (black), min. 0.00 bits (white); Local separable information: (h) positive values only, max. 29.26 bits (black), min. 0.00 bits (white), (i) negative values only, max. 0.00 bits (white), min. -18.68 bits (black).

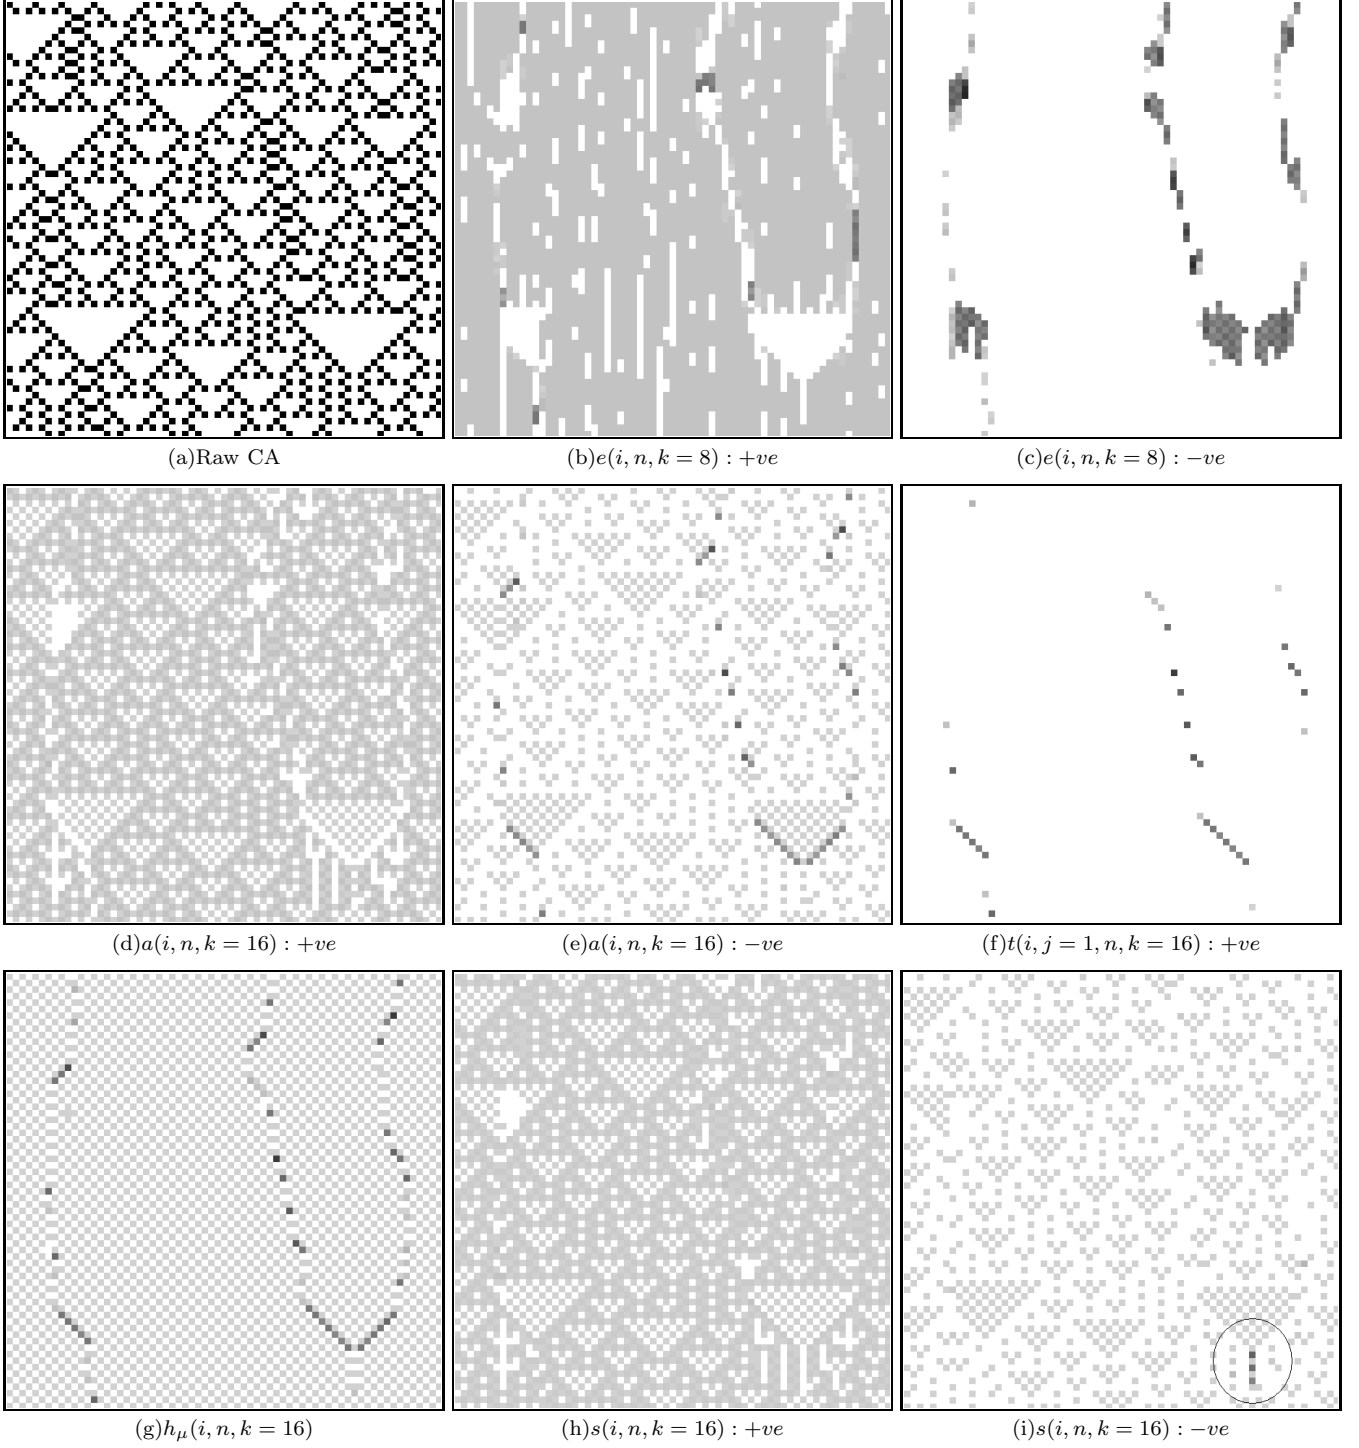


FIG. 6: Local information dynamics in rule 18: (67 time steps displayed for 67 cells): Local excess entropy: (c) negative values only, max. 0.00 bits (black), min. -8.65 bits (black); Local active information: (d) positive values only, max. 1.98 bits (black), min. 0.00 bits (white), (e) negative values only, max. 0.00 bits (white), min. -9.92 bits (black); (f) Local apparent transfer entropy (one cell to the right), positive values only, max. 11.90 bits (black), min. 0.00 bits (white); (g) Local temporal entropy rate, max. 11.90 bits (black), min. 0.00 bits (white); Local separable information: (h) positive values only, max. 1.98 bits (black), min. 0.00 bits (white), (i) negative values only, max. 0.00 bits (white), min. -14.37 bits (black).

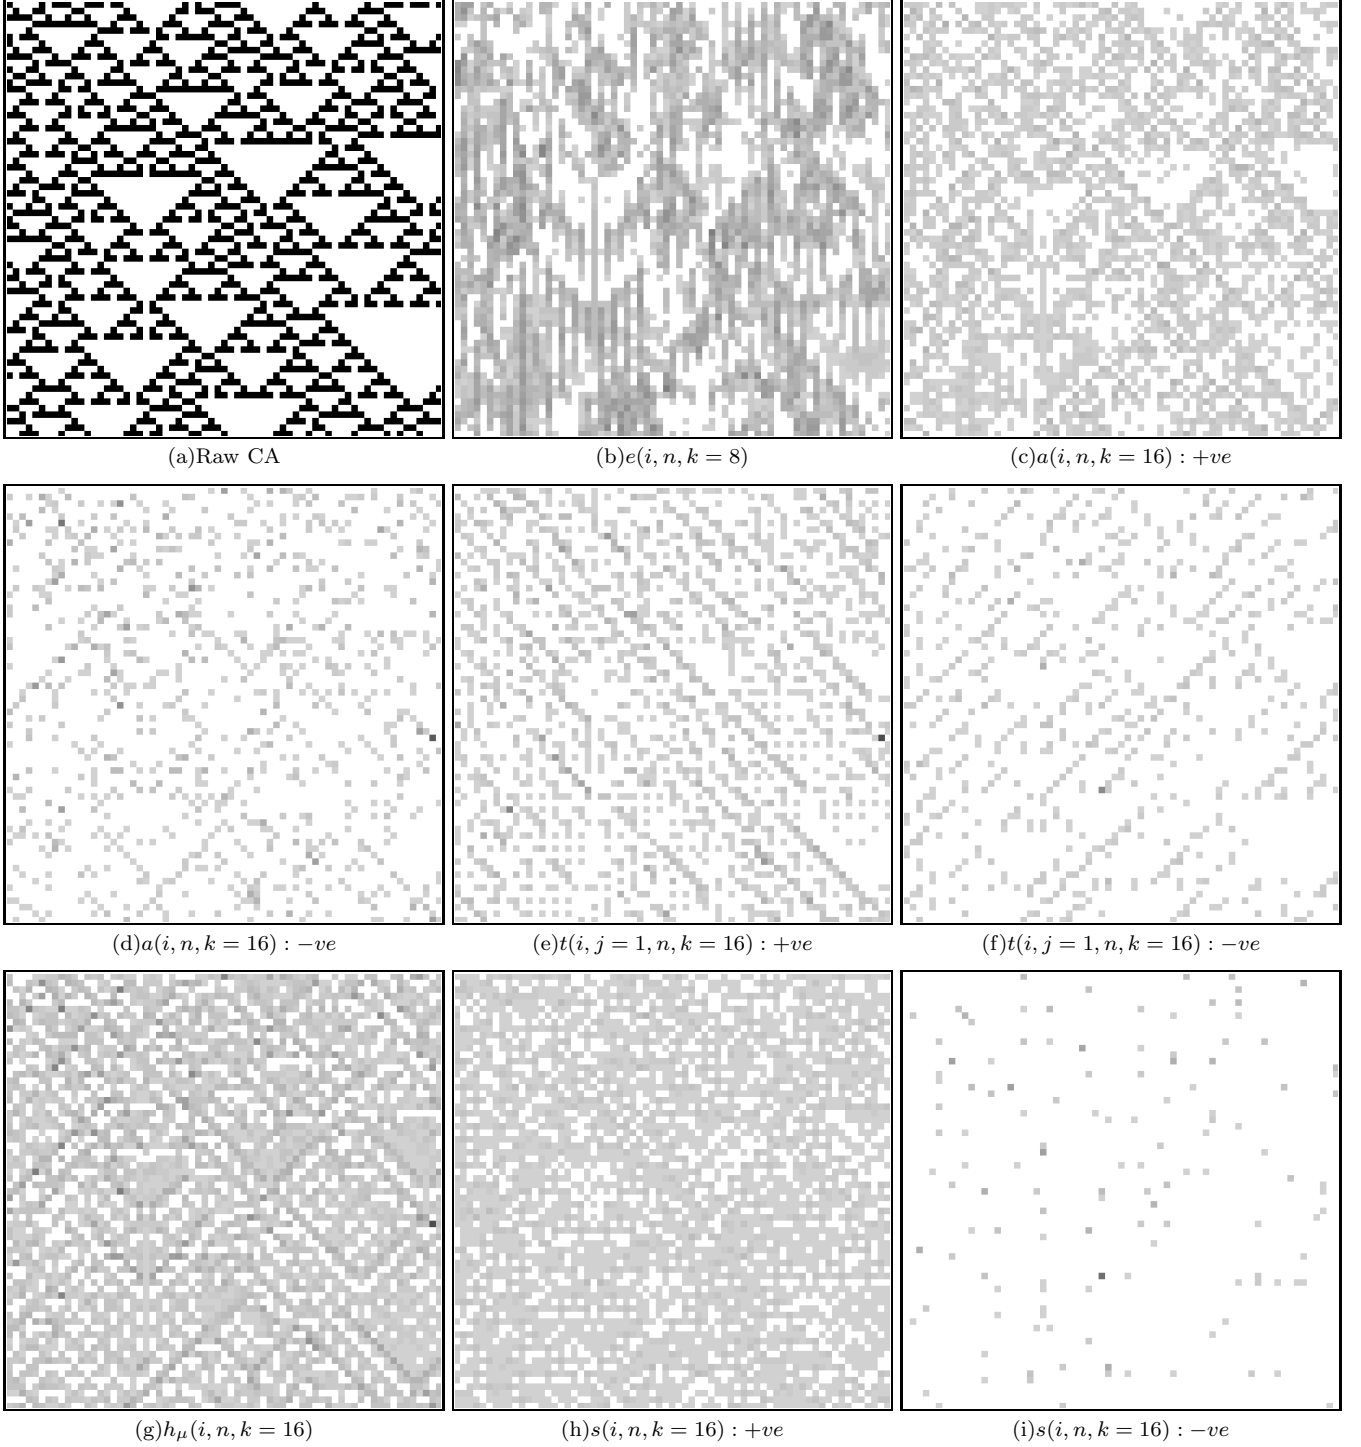


FIG. 7: Local information dynamics in rule 22: (67 time steps displayed for 67 cells): (b) Local excess entropy, positive values only, max. 4.49 bits (black), min. 0.00 bits (white); Local active information: (c) positive values only, max. 1.51 bits (black), min. 0.00 bits (white), (d) negative values only, max. 0.00 bits (white), min. -8.17 bits (black); Local apparent transfer entropy (one cell to the right): (e) positive values only, max. 9.68 bits (black), min. 0.00 bits (white), (f) negative values only, max. 0.00 bits (white), min. -7.05 bits (black); (g) Local temporal entropy rate, max. 9.68 bits (black), min. 0.00 bits (white); Local separable information: (h) positive values only, max. 5.03 bits (black), min. 0.00 bits (white), (i) negative values only, max. 0.00 bits (white), min. -14.44 bits (black).

to a *destination*. We have previously described how to measure information transfer in complex systems, and applied this to CA rules 110 and 18, in [55]. Here, we summarize this work and discuss interpretations from the perspective of computation, present more examples, and introduce collective transfer entropy.

A. Local transfer entropy

Schreiber presented *transfer entropy* as a measure for information transfer [57] in order to address deficiencies in the previous de facto measure, mutual information (Eq. (4)), the use of which he criticized in this context as a symmetric measure of statically shared information. Transfer entropy is defined as the deviation from independence (in bits) of the state transition of an information destination X from the previous state(s) of an information source Y :

$$T_{Y \rightarrow X} = \sum_{w_n} p(w_n) \log_2 \frac{p(x_{n+1}|x_n^{(k)}, y_n^{(l)})}{p(x_{n+1}|x_n^{(k)})}, \quad (20)$$

where w_n is the state transition tuple $(x_{n+1}, x_n^{(k)}, y_n^{(l)})$. As such, the transfer entropy is a *directional, dynamic* measure of information transfer. It can be viewed as a *conditional* mutual information, casting it as the average information in the source about the next state of the destination that was not already contained in the destination's past.

In [55], we demonstrated that the transfer entropy metric is an average (or *expectation value*) of a *local* transfer entropy at each observation n , i.e. $T_{Y \rightarrow X} = \langle t_{Y \rightarrow X}(n+1) \rangle$ where:

$$t_{Y \rightarrow X}(n+1) = \log_2 \frac{p(x_{n+1}|x_n^{(k)}, y_n^{(l)})}{p(x_{n+1}|x_n^{(k)})}. \quad (21)$$

For lattice systems such as CAs with spatially-ordered agents, the local information transfer to agent X_i from X_{i-j} at time $n+1$ is represented as:

$$t(i, j, n+1, k, l) = \log_2 \frac{p(x_{i,n+1}|x_{i,n}^{(k)}, x_{i-j,n}^{(l)})}{p(x_{i,n+1}|x_{i,n}^{(k)})}. \quad (22)$$

Using $l = 1$ is sensible for systems such as CAs where only the previous source state is a causal information contributor to the destination; in this case we drop l from the notation: $t(i, j, n+1, k)$. Given our focus on CAs, from here onwards we consider only $l = 1$ in this paper. This information transfer $t(i, j, n+1, k)$ to agent X_i from X_{i-j} at time $n+1$ is illustrated in Fig. 8(a). $t(i, j, n, k)$ is defined for every spatiotemporal destination (i, n) , for every information channel or direction j ; sensible values for j correspond to causal information sources, i.e. for CAs, sources within the cell range $|j| \leq r$. Again, for homogeneous agents it is appropriate to estimate the probability

distributions used in Eq. (22) from all spatiotemporal observations, and we write the average across homogeneous agents as $T(j, k) = \langle t(i, j, n, k) \rangle$.

It is important to note that the destination's own historical values can indirectly influence it via the source or other neighbors: this may be mistaken as an independent flow from the source here. In [55] we referred to this self-influence as *non-traveling information*, making analogies to standing waves. In the context of distributed computation, it is recognizable as the *active information storage*. That is, conditioning on the destination's history $x_{i,n}^{(k)}$ serves to eliminate the active information storage from the transfer entropy measurement. Yet any self-influence transmitted prior to these k values will not be eliminated: in [55] we generalized comments on the entropy rate in [57] to suggest that the asymptote $k \rightarrow \infty$ is most correct for agents displaying non-Markovian dynamics. Just as the excess entropy and active information storage require $k \rightarrow \infty$ to capture all information storage, accurate measurement of the transfer entropy requires $k \rightarrow \infty$ to eliminate all information storage from being mistaken as information transfer. The most correct form of the transfer entropy is therefore computed as:

$$t(i, j, n+1) = \lim_{k \rightarrow \infty} \log_2 \frac{p(x_{i,n+1}|x_{i,n}^{(k)}, x_{i-j,n})}{p(x_{i,n+1}|x_{i,n}^{(k)})}, \quad (23)$$

with $t(i, j, n+1, k)$ retained for finite- k estimates.

Additionally, the transfer entropy may be conditioned on other possible causal information sources, to eliminate their influence from being attributed to the source in question Y [57]. In general, this means conditioning on all sources Z in the set of causal information contributors V (except for Y) with joint state $v_{y,n}$, giving the local *complete* transfer entropy [55]:

$$t_{Y \rightarrow X}^c(n+1, k) = \log_2 \frac{p(x_{n+1}|x_n^{(k)}, y_n, v_{y,n})}{p(x_{n+1}|x_n^{(k)}, v_{y,n})}, \quad (24)$$

$$v_{y,n} = \{z_n | \forall Z \in V, Z \neq Y, X\}. \quad (25)$$

For ECAs this means conditioning on other sources $v_{i,j,n}^r$ in the neighborhood of the destination to obtain [55]:

$$t^c(i, j, n+1, k) = \log_2 \frac{p(x_{i,n+1}|x_{i,n}^{(k)}, x_{i-j,n}, v_{i,j,n}^r)}{p(x_{i,n+1}|x_{i,n}^{(k)}, v_{i,j,n}^r)}, \quad (26)$$

$$v_{i,j,n}^r = \{x_{i+q,n} | \forall q : -r \leq q \leq +r, q \neq -j, 0\}. \quad (27)$$

Again, the most correct form is $t^c(i, j, n+1)$ in the limit $k \rightarrow \infty$. In deterministic systems (e.g. CAs), complete conditioning renders $t^c(i, j, n) \geq 0$ because the source can only add information about the outcome of the destination. Calculations conditioned on no other information contributors (as in Eq. (22)) are labeled as *apparent* transfer entropy. Local apparent transfer entropy $t(i, j, n)$ may be either positive or negative, with negative values occurring where (given the destination's history) the source element is actually misleading about the next state of the destination.

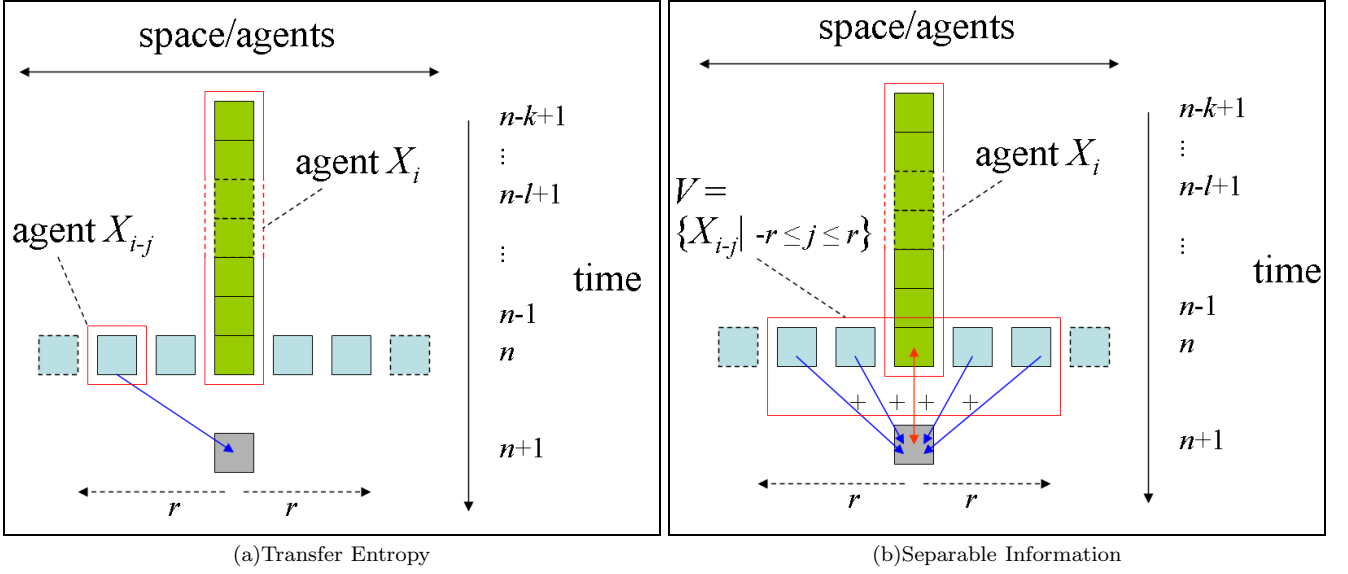


FIG. 8: (a) Transfer Entropy $t(i, j, n+1, k)$: information contained in the source cell X_{i-j} about the next state of the destination cell X_i at time $n+1$ that was not contained in the destination's past. (b) Separable information $s(i, n+1, k)$: information gained about the next state of the destination from separately examining each causal information source in the context of the destination's past. For ECAs these causal sources are within the cell range r .

B. Total information, entropy rate and collective information transfer

The total information required to predict the next state of any agent i is the *local entropy* $h(i, n+1)$, where the entropy is the average of these local values: $H(X_i) = \langle h(i, n+1) \rangle$. Similarly, the *local temporal entropy rate* $h_\mu(i, n+1, k)$ is the information to predict the next state of agent i given that agent's past, and the entropy rate is the average of these local values: $H_\mu(X_i, k) = \langle h_\mu(i, n+1, k) \rangle$. As demonstrated in Appendix A, the local entropy can be considered as the sum of the local active information storage $a(i, n+1, k)$ and local temporal entropy rate:

$$h(i, n+1) = a(i, n+1, k) + h_\mu(i, n+1, k). \quad (28)$$

For deterministic systems (e.g. CAs) there is no intrinsic uncertainty, so the local temporal entropy rate is equal to the *local collective transfer entropy* (see Appendix A) and represents a collective information transfer: the information about the next state of the destination *jointly* added by the causal information sources that was not contained in the past of the destination. Also, Appendix A shows that (via a sum of incrementally conditioned mutual information terms) for ECAs we have:

$$h(i, n+1) = a(i, n+1, k) + t(i, j = -1, n+1, k) + t^c(i, j = 1, n+1, k), \quad (29)$$

(and vice-versa in $j = 1, -1$).

Clearly, this total information is not simply a simple sum of the active information transfer and the apparent transfer entropy from each source, nor the sum of

the active information transfer and the complete transfer entropy from each source. In earlier work [55], we demonstrated that the sum of transfer entropies from each source (either apparent or complete) formed a useful single spatiotemporal filter for emergent structure in CAs, whereas the transfer entropy from each source displays information transfer in one given direction only. Given Eq. (28), we suggest that the local collective transfer entropy (or simply the local temporal entropy rate $h_\mu(i, n, k)$ for deterministic systems) is likely to be a more meaningful measure and filter for this purpose.

C. Local information transfer results

Local complete and apparent transfer entropy were applied to ECA rules 110 and 18 in [55]. Here we revisit these and present further examples, focusing on the local apparent transfer entropy: profiles of the positive values of $t(i, j = 1, n, k = 16)$ are plotted for rules 54 (Fig. 2(e)), ϕ_{par} (Fig. 5(e)), 18 (Fig. 6(f)) and 22 (Fig. 7(e)), with $t(i, j = -1, n, k = 16)$ plotted for rule 110 (Fig. 4(e)). We also measure the profiles of the local temporal entropy rate $h_\mu(i, n+1, k)$ (which is equal to the local collective transfer entropy in these deterministic systems) here in Fig. 4(g) for rule 110, Fig. 5(g) for ϕ_{par} , and Fig. 6(g) for rule 18.

Both the local apparent and complete transfer entropy highlight particles (including gliders and domain walls) as strong positive information transfer against background domains. Importantly, the particles are measured as information transfer in their direction of macroscopic motion, as expected. For example, at the "x"

marks in Fig. 3 which denote parts of the right-moving γ^+ gliders, we have $p(x_{i,n+1}|x_{i,n}^{(k=16)}, x_{i-1,n}) = 1.00$ and $p(x_{i,n+1}|x_{i,n}^{(k=16)}) = 0.25$: there is a strong information transfer of $t(i, j = 1, n, k = 16) = 2.02$ bits here because the source (in the glider) added a significant amount of information to the destination about the continuation of the glider. For ϕ_{par} we confirm the role of the gliders as information transfer agents in the human understandable computation, and demonstrate information transfer across multiple units of space per unit time step for fast-moving gliders in Fig. 5(f). Interestingly, we also see in Fig. 5(e) and Fig. 5(f) that the apparent transfer entropy can attribute information transfer to several information sources, whereas the complete transfer entropy (not shown) is more likely to attribute the transfer to the single causal source (though we emphasize that information transfer and causality are distinct concepts).

As expected, the local temporal entropy rate profiles $h_\mu(i, n+1, k)$ highlight particles moving in each relevant channel and are a useful single spatiotemporal filter for emergent structure. In fact, these profiles are quite similar to the profiles of the negative values of local active information. This is not surprising given they are counterparts in Eq. (28): where $h_\mu(i, n+1, k)$ is strongly positive (i.e. greater than 1 bit), it is likely that $a(i, n+1, k)$ is negative since the local single cell entropy will average close to 1 bit for these examples. Unlike $a(i, n+1, k)$ however, the local temporal entropy rate $h_\mu(i, n+1, k)$ is never negative.

In [55] we also demonstrated that while there is zero information transfer in an infinite periodic domain, there is a small non-zero information transfer in domains acting as a background to gliders, effectively indicating the absence of gliders. These small non-zero information transfers are stronger in the wake of a glider, indicating the absence of (relatively common) following gliders. Similarly, we note here that the local temporal entropy rate profiles $h_\mu(i, n+1, k)$ contain small but non-zero values in these periodic domains. Furthermore, there is interesting structure to the information transfer in the domain of rule 18 (see Fig. 6(g)). As described in Section IV D, this domain is of spatial and temporal period 2, with every second site being “0” and every other site being either a “0” or a “1”. Since the “0”’s at every second site are completely predictable (in the absence of domain walls) given their past history, $h_\mu(i, n+1, k)$ at these points approaches zero bits. On the other hand, at every other site $h_\mu(i, n+1, k)$ approaches 1 bit since observations of a “0” or a “1” are roughly equally likely with the past history indicating this alternate phase of the background. This result complements our observations in [55] of the local transfer entropies here. As shown for $t(i, j = 1, n+1, k = 16)$ in Fig. 6(f), local apparent transfer entropy approaches 0 bits at each site in the domain: the “0”’s at every second site are completely predictable from their pasts, while the alternate sites require both neighboring sources to predict the outcome. The local

complete transfer entropies therefore measure approximately 1 bit in their ability to finally determine the sites in the alternate phase, but measure no transfer for the “0” sites in the primary phase. Summing the respective pairs of apparent and complete transfer entropies (as per Eq. (29)) therefore matches our results for $h_\mu(i, n+1, k)$ as required. With these results, local transfer entropy provided the first quantitative evidence for the long-held conjecture that particles are the *dominant* (but not the only) information transfer agents in CAs.

The highlighting of structure by local transfer entropy is similar to results from other methods of filtering for structure in CAs [21, 22, 31, 32], but subtly different in revealing the leading edges of gliders as the major information transfer elements in the glider structure, and providing multiple profiles (one for each direction or channel of information transfer). Note that while achieving the limit $k \rightarrow \infty$ is not computationally infeasible, at least a significant k was required to achieve a reasonable estimates of the transfer entropy; without this, the active information storage was not eliminated from the transfer entropy measurements in the domains, and the measure did not distinguish the particles from the domains [55].

Also, a particularly relevant result for our purposes is the finding of negative values of transfer entropy for some space-time points in particles moving orthogonal to the direction of measurement. This is displayed for $t(i, j = 1, n, k = 16)$ in rule 54 (Fig. 2(f)), and $t(i, j = -1, n, k = 16)$ for rule 110 (Fig. 4(f)), and also occurs for rule 18 (results not shown). In general this is because the source, as part of the domain, suggests that this same domain found in the past of the destination will continue; however since the next state of the destination forms part of the particle, this suggestion proves to be misinformative. For example, consider the “x” marks in Fig. 3 which denote parts of the right-moving γ^+ gliders. If we now examine the source at the right (still in the domain), we have $p(x_{i,n+1}|x_{i,n}^{(k=16)}, x_{i+1,n}) = 0.13$ giving $t(i, j = 1, n, k = 16) = -0.90$ bits: this is negative because the source (still in the domain) was misinformative about the destination.

Regarding the local information transfer structure of rule 22, we note similar results as for local information storage. There is much information transfer here (in fact the average value $T(j = 1, k = 16) = 0.19$ bits is greater than for rule 110 at 0.07 bits), although there is no coherent structure to this transfer. Again, this demonstrates the utility of local information metrics in providing more detailed insights into system dynamics than their global averages.

In this section, we have described how the local transfer entropy quantifies the information transfer at space-time points within a system, and provides evidence that particles are the dominant information transfer agents in CAs. We have also introduced the collective transfer entropy to quantify the joint information contribution from all causal information contributors, and measured this in deterministic systems using the temporal entropy rate.

However, we have not yet separately identified collision events in CAs: to complete our exploration of the information dynamics of computation, we now consider the nature of information modification.

VI. INFORMATION MODIFICATION

Langton interpreted information modification as interactions between transmitted and/or stored information which resulted in a modification of one or the other [3]. CAs provide an illustrative example, where the term *interactions* is generally interpreted to mean collisions of particles (including blinkers as information storage), with the resulting dynamics involving something other than the incoming particles continuing unperturbed. The resulting dynamics could involve zero or more particles (with an annihilation leaving only a background domain), and perhaps even some of the incoming particles. The number of particles resulting from a collision has been studied elsewhere [43]. Given the focus on perturbations in the definition here, it is logical to associate a collision event with the modification of transmitted and/or stored information, and to see it as an information processing or decision event. Indeed, as an information processing event the important role of collisions in determining the dynamics of the system is widely acknowledged [43], e.g. in the ϕ_{par} density classification.

Attempts have previously been made to quantify information modification or processing in a system [14, 16, 17]. However, these have either been too specific to allow portability across system types (e.g. by focusing on the capability of a system to solve a known problem, or measuring properties related to the particular type of system being examined), focus on general processing as movement or interpretation of information rather than specifically the modification of information, or are not amenable to measuring information modification at *local* space-time points *within* a distributed system. In this section, we present the separable information as a tool to detect non-trivial information modification events, and demonstrate it as the first measure to identify collisions in CAs as such.

A. Local separable information

We begin by considering what it means for a particle to be *modified*. For the simple case of a glider, a modification is simply an alteration to the predictable periodic pattern of the glider's dynamics. At such points, an observer would be surprised or misinformed about the next state of the glider, having not taken account of the entity about to perturb it. This interpretation is a clear reminder of our earlier comments that local apparent transfer entropy $t(i, j, n)$ and local active information storage $a(i, n)$ were negative where the respective information sources were *misinformative* about the next state of the

information destination (in the context of the destination's past for transfer entropy). Local active information storage was misinformative at gliders, and local apparent transfer entropy was misinformative at gliders traveling in the orthogonal direction to the measurement. This being said, one expects that the local apparent transfer entropy measured in the direction of glider motion will be *more informative* about its evolution than any misinformation conveyed from other sources. However, where the glider is modified by a collision with another glider, we can no longer expect the local apparent transfer entropy in its macroscopic direction of motion to remain informative about its evolution. Assuming that the incident glider is also be perturbed, the local apparent transfer entropy in its macroscopic direction of motion will also not be informative about its evolution at this collision point. We expect the same argument to be true for irregular particles, or domain walls.

As such, we make the hypothesis that at the spatiotemporal location of a local information modification event or collision, *separate* inspection of each information source will *misinform* an observer overall about the next state of the modified information destination. More specifically, the information sources referred to here are the past history of the destination (via the local active information storage) and each other causal information contributor (examined in the context of the past history of the destination, via their local apparent transfer entropies).

We quantify the total information gained from separate observation of the information storage and information transfer contributors as the *local separable information* $s_X(n)$:

$$s_X(n) = a_X(n) + \sum_{Y \in V, Y \neq X} t_{Y \rightarrow X}(n), \quad (30)$$

with the subscripts indicating the destination and source variables. Again, the separable information S_X denotes the average $S_X = \langle s_X(n) \rangle$. For CAs, where the causal information contributors are homogeneously within the neighborhood r , we write the local separable information in lattice notation as:

$$s(i, n) = a(i, n) + \sum_{j=-r, j \neq 0}^{+r} t(i, j, n). \quad (31)$$

We use $s(i, n, k)$ to represent finite- k estimates, and show $s(i, n, k)$ diagrammatically in Fig. 8(b).[73]

As inferred earlier, we expect the local separable information to be *positive* or *highly separable* where separate observations of the information contributors are informative overall regarding the next state of the destination. This may be interpreted as a trivial information modification, because information storage and transfer are not interacting in any significant manner. More importantly, we expect the local separable information to be *negative* at spatiotemporal points where an information modification event or collision takes place. Here, separate observations are misleading overall because a *non-trivial information modification* is taking place (i.e. the

information storage and transfer are interacting. It is thus clear how we can understand information modification as the interaction between information storage and information transfer.

Importantly, this formulation of non-trivial information modification aligns with the descriptions of complex systems as consisting of (a large number of) elements interacting in a *non-trivial* fashion [24], and of emergence as where “*the whole is greater than the sum of its parts*”. Here, we quantify the sum of the parts in $s(i, n)$, and “the whole” refers to examining all information sources together; the whole is greater where all information sources must be examined together in order to receive positive information on the next state of the examined entity. That being said, there is no quantity representing “the whole” as such, simply the indication that the sources must be examined *together*. We emphasize that $s(i, n)$ is not the total information an observer needs to predict the state of the destination; this is measured by the single-site entropy $h(i, n)$ (see Section A).

Finally, we introduce the notation $S^+(k)$ and $S^-(k)$ as the averages of positive and negative local values of $s(i, n, k)$ in contributing to the average $S(k)$. We have for example $S^+(k) = \langle s^+(i, n, k) \rangle$, where:

$$s^+(i, n, k) = \begin{cases} s(i, n, k) & \text{if } s(i, n, k) \geq 0 \\ 0 & \text{if } s(i, n, k) < 0 \end{cases}, \quad (32)$$

while $S^-(k) = \langle s^-(i, n, k) \rangle$ is defined in the opposite manner.

B. Local separable information results

The simple gliders in ECA rule 54 give rise to relatively simple collisions which we focus on in our discussion here. The positive values of $s(i, n, k = 16)$ for rule 54 are displayed in Fig. 2(g); notice that these are concentrated in the domain regions and at the stationary gliders (α and β). As expected, these regions are undertaking trivial computations only. Fig. 2(h) displays the negative values of $s(i, n, k = 16)$, with their positions marked in Fig. 2(i). The dominant negative values are clearly concentrated around the areas of collisions between the gliders, including collisions between the traveling gliders only (marked by “A”) and between the traveling gliders and the stationary gliders (marked by “B”, “C” and “D”).

Collision “A” involves the γ^+ and γ^- particles interacting to produce a β particle ($\gamma^+ + \gamma^- \rightarrow \beta$ [43]). The only information modification point highlighted is one time step below the point at which the gliders naively appear to collide (see close-up of raw states in Fig. 3). The periodic pattern in the past of the destination breaks there, however the neighboring sources are still able to support separate prediction of the state (i.e. $a(i, n, k = 16) = -1.09$ bits, $t(i, j = 1, n, k = 16) = 2.02$ bits and $t(i, j = -1, n, k = 16) = 2.02$ bits, giving $s(i, n, k = 16) = 2.95$ bits). This is no longer the case

however where our metric has successfully identified the modification point; there we have $a(i, n, k = 16) = -3.00$ bits, $t(i, j = 1, n, k = 16) = 0.91$ bits and $t(i, j = -1, n, k = 16) = 0.90$ bits, with $s(i, n, k = 16) = -1.19$ bits suggesting a non-trivial information modification. A delay is also observed before the identified information modification points of collision types “B” ($\gamma^+ + \beta \rightarrow \gamma^-$, or vice-versa in γ -types), “C” ($\gamma^- + \alpha \rightarrow \gamma^- + \alpha + 2\gamma^+$, or vice-versa) and “D” ($2\gamma^+ + \alpha + 2\gamma^- \rightarrow \alpha$); possibly these delays represent a time-lag of information processing. Not surprisingly, the results for these other collision types imply that the information modification points are associated with the creation of new behavior: in “B” and “C” these occur along the newly created γ gliders, and for “C” and “D” in the new α blinkers.

Importantly, weaker information modification points continue to be identified at every second point along all the γ^+ and γ^- particles after the initial collisions (these are too weak to appear in Fig. 2(h) but can be seen for a similar glider in rule 110 in Fig. 4(i)). This was unexpected from our earlier hypothesis. However, these events can be understood as non-trivial computations of the continuation of the glider in the *absence* of a collision; in effect they are virtual collisions between the real glider and the absence of an incident glider. These weak collision events are more significant in the wake of real collisions, since incident gliders are relatively more likely in these areas. Interestingly, this finding is analogous to the small but non-zero information transfer in periodic domains indicating the absence of gliders.

We also note that measurements of local independently observed information must be performed with a reasonably large value of k . Earlier, we observed that for appropriate measurement of information storage and transfer k should be selected to be as large as possible for accuracy, at least larger than the scale of the period of the regular background domain for filtering purposes. Here, using $k < 4$ could not distinguish any information modification points clearly from the domains and particles, and even $k < 8$ could not distinguish *all* the modification points (results not shown). Correct quantification of information modification requires satisfactory estimates of information storage and transfer, and accurate distinction between the two.

We observe similar results in the profile of $s(i, n, k = 10)$ for ϕ_{par} , confirming that the particle collisions here are non-trivial information modification events, and therefore completing the evidence for all of the conjectures about this human understandable computation.

The results for $s(i, n, k = 16)$ for ECA rule 110 (see Fig. 4(h) and Fig. 4(i)) are also similar to those for rule 54. Here, we have collisions “A” and “B” which show non-trivial information modification points slightly delayed from the collision in a similar fashion to those for rule 54. We note that collisions between some of the more complex glider structures in rule 110 (not shown) exhibit non-trivial information modification points which are more difficult to interpret, and which are even more

delayed from the initiation of the collision. The larger delay is perhaps this is a reflection of the more complex gliders requiring more time steps for the processing to take place. An interesting result not seen for rule 54 is a collision where an incident glider is absorbed by a blinker (see label “C”), without any modification to the absorbing glider. No information modification is detected for this absorption event by $s(i, n, k = 16)$: this is as expected because the information storage for the absorbing blinker is sufficient to predict the dynamics at this interaction.

As a further test of the measure, we examine collisions between the domain walls of rule 18. As displayed in Fig. 6(i), the collision between the domain walls is quite clearly highlighted as the dominant information modification event for this rule. The initial information modification event is clearly where one would naively identify the collision point, yet it is followed by two secondary information modification points separated by two time steps. At the raw states of these three collision points in Fig. 6(a), the outer domains have effectively coalesced, however the observer cannot be certain that the new domain has taken hold at this particular cell until observing a “1” at the alternate phase. As such, information modification events are observed at each point in the new alternate phase until a “1” confirms the outer domains have joined. This could provide a parallel to the observation of delays in information processing observed earlier. Importantly, this result provides evidence that collision of irregular particles are information modification events, as expected. It is also worth noting that these collisions always result in the destruction of the domain walls (and the inner domain), indicating that our method captures destruction events as well as creation. (This is also true for the $\gamma^+ + \gamma^- + \beta \rightarrow \emptyset$ event in rule 54, not shown). Also, as displayed in Fig. 6(h) and Fig. 6(i), the background domain takes values of $s(i, n, k = 16)$ as either positive or negative with $a(i, n, k = 16)$, since $t(i, j = 1, n, k = 16)$ and $t(i, j = -1, n, k = 16)$ vanish at these points. This indicates that some minor information processing is required to compute the “0” sites in the alternate phase (whereas the “0” sites for every second point and the “1”s in the alternate phase are trivial computations dominated by information storage). Finally, the domain walls here appear to give rise to only positive values of $s(i, n, k = 16)$. This indicates that the domains walls contain only trivial information modification, in contrast with the gliders in rule 54 which required a small amount of non-trivial information processing in order to compute their continuation. This is perhaps akin to the observation in [21] that the domain walls in rule 146 are largely determined by the dynamics on either side, i.e. they are not the result of any interaction per se but of dominance from a single source at each time step.

We also apply $s(i, n, k = 16)$ to ECA rule 22, as displayed in Fig. 7(i) and Fig. 7(j). As could be expected from our earlier results, there are many points of both positive and negative local separable information here.

The presence of negative values implies the occurrence of non-trivial information modification, yet there does not appear to be any structure to these profiles. Again, this aligns well with the lack of coherent structure found using the other measures in this framework and from the local statistical complexity profile of rule 22 [21].

Here, we have introduced the local separable information to quantify information modification at each spatiotemporal point in a complex system. Information modification events occur where the separable information is negative, indicating that separate or independent inspection of the causal information sources (in the context of the destination’s past) is misleading because of non-trivial interaction between these sources. The local separable information was demonstrated to provide the first quantitative evidence that particle collisions in CAs are the dominant information modification events therein. The measure is capable of identifying events involving both creation and destruction, and interestingly the location of an information modification event often appears delayed perhaps due to a time-lag in information processing.

VII. IMPORTANCE OF COHERENT COMPUTATION

Our framework has proven successful in locally identifying the component operations of distributed computation. We now consider whether this framework can provide any insights into the overall complexity of computation. In other words, what can our results say about the difference in the complex computations of rules 110 and 54 as compared to rule 22 and others?

We observe that the *coherence* of local computational structure appears to be the most significant differentiator here. “Coherence” implies a property of sticking together or a logical relationship [58]: in this context we use the term to describe a logical spatiotemporal relationship between values in local information dynamics profiles. For example, the manner in which particles give rise to similar values of local transfer entropy amongst spatiotemporal neighbors is coherent. From the spatiotemporal profiles presented here, we note that rules 54 and 110 exhibit the largest amount of coherent computational structure, with rule 18 containing a smaller amount of less coherent structure. Rules 22 and 30 (results for rule 30 not shown) certainly exhibit all of the elementary functions of computation, but do not appear to contain any coherent structure to their computations. This aligns well with similar explorations of local information structure for these rules, e.g. [21]. Using language reminiscent of Langton’s analysis [3], we suggest that complex systems exhibit very *highly-structured coherent* computation in comparison to ordered systems (which exhibit coherence but minimal structure in a computation dominated by information storage) and chaotic systems (whose computations are dominated by rampant information transfer

eroding any coherence).

The key question then is whether one can perform any meta-analysis on the local information dynamics in order to quantify the differences in complexity of these computations in terms of coherent structure. It is unlikely that such meta-analyses will produce system-wide measures of the complexity of computation (see comments in Section III B), however we attempt to at least produce useful heuristics for this purpose. Here, we present three candidate approaches: examining the *average* information dynamics, correlation analyses and state-space plots of the local values.

An obvious first step is to check whether the average information dynamics provide useful summaries regarding the coherence and complexity of computation in each CA rule (despite the fact that the local values themselves provide far more detail). These averages are presented in Table I. One striking feature of the known complex rules is that the apparent transfer entropy in *each* channel is a large proportion of the complete transfer entropy for that channel. Apparent transfer entropy can only be high where the source has a clear coherent influence on the destination, while complete transfer entropy can separately be high where the influence of the source is mediated through an interaction. In the complex CAs, single sources can influence destinations without needing to interact with other sources, supporting the movement of coherent particle structures. Importantly, this occurs for multiple channels, meaning that we have bidirectional traveling coherent structures that should *interact* at some point. A similar feature is that their separable information approaches the entropy (again indicating dominance of single sources), along with a very low proportion of non-trivial information modification events (indicated by an almost vanishing S^- and a small proportion of points with $s(i, n) < 0$). Given our knowledge of the importance of these events to computation, their shortage in complex computation initially seems counter-intuitive. However, we suggest that the power of these events lies in their subtlety: used judiciously they allow a complex *coherent* computation, but occurring too often they disturb the coherence of the computation which then becomes chaotic. Similarly, we note that chaotic rules exhibit higher values of the complete transfer entropy along with lower values of the apparent transfer entropy; this provides another indication of significant interaction between components eroding the coherent computation in this regime. While these observations quantify neither coherence nor complexity of computation, they do provide useful heuristics for identifying those properties.

Our interpretation of coherence as meaning a logical spatiotemporal relationship between local values suggests that it may be measured via the autocorrelation within profiles of each of their local information dynamics. Table II shows for example the autocorrelation for local transfer entropy values $t(i, j = 1, n, k = 16)$ separated by 1 step in time and 1 step to the right in space. The separation for the autocorrelation here is the same as the

interval across which the local transfer entropy is measured. Notice that the rules exhibiting particles (110, 54 and 18) display the highest correlation values here, since coherent particles exhibit spatiotemporal correlation in the direction of particle motion. Similar results are observed for $t(i, j = -1, n, k = 16)$ with autocorrelation over 1 step in time and 1 step to the left in space. Again, this observation is a useful heuristic, and parallels the above observations regarding the proportion of average transfer entropy values.

Coherence may also be interpreted as a logical relationship *between* profiles of the individual local information dynamics (as three axes of complexity) rather than only within them. To investigate this possibility, Fig. 9 plots state-space diagrams of the local apparent transfer entropy for $j = 1$ versus local active information storage, while Fig. 10 plots the local separable information versus local active information storage for several CA rules. Each point in these diagrams represents the local values of each measure at one spatiotemporal point, thereby generating a complete state-space for the CA. Such state-space diagrams are known to provide insights into structure that are not visible when examining either metric in isolation; for example, in examining structure in classes of systems (such as logistic maps) by plotting average excess entropy versus entropy rate while changing a system parameter [59]. Here however we are looking at structure *within* a single system rather than across a class of systems.

The state-space diagrams for rule 110 (Fig. 9(a) and Fig. 10(a)) exhibit interesting structure, with significant clustering around certain areas and lines in the state space, reflecting its status as a complex rule. (The two diagonal lines are upper limits representing the boundary condition $t^c(i, j = -1, n, k = 16) \geq 0$ for both destination states “0” and “1”). On the other hand, the example state space diagrams for rule 30 (Fig. 9(b) and Fig. 10(b)) exhibit minimal structure (apart from the mathematical upper limit), with a smooth spread of points across the space reflecting its underlying chaotic nature. From the apparent absence of coherent structure in its space-time information profiles, one may expect state-space diagrams for rule 22 to exhibit a similar absence of structure to rule 30. As shown by Fig. 9(c) and Fig. 10(c) however this is not the case: the state-space diagrams for rule 22 exhibit significant structure, with similar clustering to that of rule 110.

To attempt to quantify the coherence of the structure here, we measure the correlation coefficient between the values of $t(i, j = 1, n, k = 16)$ and $a(i, n, k = 16)$ for example (see Table II). However, the correlation measures linear relationships alone; we also measure the mutual information[74] (see Table II) between the pairs $(t(i, j = 1, n, k = 16), a(i, n, k = 16))$ and $(s(i, n, k = 16), a(i, n, k = 16))$ as a more general measure of their underlying relationship. The mutual information results suggest (as expected) that rules 110 and 54 display strong relationships between all measures, with signifi-

TABLE I: Table of average information dynamics (all with $k = 16$, values to 2 decimal places except S^- for rule 110) with units in bits (except for $p(s(i, n) < 0)$, the proportion of space-time points with negative local separable information), for several ECA rules.

ECA rule	H	H_μ	A	$T_{j=1}$	$T_{j=-1}$	$T_{j=1}^c$	$T_{j=-1}^c$	S	S^+	S^-	$p(s(i, n) < 0)$
110	0.99	0.18	0.81	0.07	0.11	0.07	0.11	0.98	0.98	-0.002	0.003
54	1.00	0.27	0.73	0.08	0.08	0.19	0.19	0.89	0.90	-0.01	0.03
22	0.93	0.75	0.19	0.19	0.19	0.56	0.56	0.56	0.62	-0.05	0.09
18	0.82	0.53	0.29	0.01	0.01	0.52	0.52	0.32	0.46	-0.14	0.25
30	1.00	0.99	0.01	0.73	0.01	0.98	0.26	0.75	0.82	-0.07	0.08

TABLE II: Table of the autocorrelation between values of local transfer entropy separated by 1 step in time and space ($t = t(i, j = 1, n, k = 16)$, $t' = t(i + 1, j = 1, n + 1, k = 16)$), correlation coefficient and mutual information (in bits) between values of local active information storage ($a = a(i, n, k = 16)$) and local transfer entropy (t) at the same space-time points, and mutual information (in bits) between values of local separable information ($s = s(i, n, k = 16)$) and local active information storage (a) for several ECA rules.

ECA rule	$C_{tt'}$	C_{ta}	$I_{t;a}$	$I_{s;a}$
110	0.19	-0.57	0.35	0.69
54	0.45	-0.54	0.58	0.57
22	0.09	-0.28	0.25	0.40
18	0.44	-0.23	0.09	1.48
30	0.03	-0.19	0.17	0.00

cant strength in these relationships for rule 22 also, and a strong relationship for $(s(i, n, k = 16), a(i, n, k = 16))$ in rule 18 (as these are highly correlated in the period-2 pattern in its domain). Again, this appears to be a useful heuristic for coherence in computation, though the state-space diagrams themselves contains much more detail about the relationships between these axes of complexity.


Importantly, the apparent information structure in the state-space diagrams lends some credence to the claims of complex behavior for rule 22 discussed in Section III C. However it is a very subtle type of structure, not complex enough to be revealed in the individual local information profiles shown here or by other authors (e.g. [21]). The structure does not appear to be coherent in these individual profiles, though the state space diagrams indicate a coherent relationship between the local information dynamics which may underpin coherent computation at other scales. Given the subtlety of structure in the bounds of our analysis, and using our mutual information heuristics, at this stage we conclude that the behavior of this rule is less complex than that exhibited by rules 110 and 54.

Here we have suggested that coherent information structure is a defining feature of complex computation, and presented a number of important techniques and heuristics for inferring this property using local information dynamics. A particular example are state-space diagrams for local information dynamics, which produce useful visual results and were shown to provide interesting insight into the nature of computation in rule 22.


VIII. CONCLUSION

We have presented a complete quantitative framework for the information dynamics of distributed computation in complex systems. Our framework quantifies the information dynamics in terms of the component operations of universal computation: information storage, information transfer and information modification. Importantly, the framework describes the manner in which information storage and transfer interact to produce non-trivial computation where “the whole is greater than the sum of the parts”. Our framework places particular importance on examining computation on a local scale. While averaged or system-wide measures have their place in providing summarized results, this focus on the local scale is vital for understanding the information dynamics of computation and provides many insights that averaged measures cannot.


We applied the framework to cellular automata, an important example because of the weight of previous studies on the nature of distributed computation in these systems. Significantly, our framework provides quantitative evidence for the widely accepted conjectures that blinkers provide information storage in CAs, particles are the dominant information transfer agents, and particle collisions are the dominant information modification events. In particular, this was demonstrated for the human-understandable density classification computation carried out by the rule ϕ_{par} . This is a fundamental contribution to our understanding of the nature of distributed computation, and provides impetus for the framework to be used for the analysis and design of other complex systems.



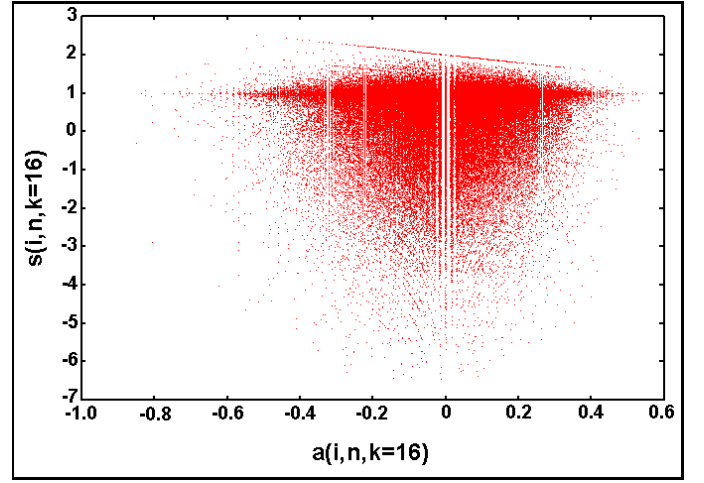
(a)110




(a)110




(b)30



(b)30



(c)22



(c)22

FIG. 9: State space diagrams of local transfer entropy (one step to the right) $t(i, j = 1, n, k = 16)$ versus local active information $a(i, n, k = 16)$ at the same space-time point (i, n) for several ECA rules: (a) 110, (b) 30 and (c) 22.

FIG. 10: State space diagrams of local separable information $s(i, n, k = 16)$ versus local active information $a(i, n, k = 16)$ at the same space-time point (i, n) for several ECA rules: (a) 110, (b) 30 and (c) 22.

The application to CAs aligned well with other methods of filtering for complex structure in CAs. However, our work is distinct in that it provides several different views of the system corresponding to each type of computational structure. In particular, the results align well with the insights of computational mechanics, providing a strong connection between this field and the local information dynamics of computation.

From our results, we also observed that coherent local information structure is a defining feature of complex distributed computation, and presented a number of techniques to meta-analyze local information dynamics in order to infer coherent complex computation. Here, our framework provides further insight into the nature of computation in rule 22 with respect to the accepted complex rules 54 and 110. Certainly rule 22 exhibits all of the elementary functions of computation, yet (in line with [21]) there is no apparent coherent structure to the profiles of its local information dynamics. On the other hand, state space views of the interplay between these local information dynamics reveal otherwise hidden structure. Our framework is unique in its ability to resolve both of these aspects. We conclude that rule 22 exhibits more structure than chaotic rules, yet the subtlety of this structure prevent it from being considered as complex than rules 110 and 54.

The major thrust of our future work is to apply this framework to other systems, e.g. we are examining computation in random Boolean networks as models of gene regulatory networks [60] and in modular robotics [61]. Given the information-theoretic basis of this framework, it is readily applicable to other systems. Furthermore, we intend to explore the relationship between the information dynamics description of distributed computation and other perspectives of computation, for example descriptions of collective computation (e.g. in the light-cone model of computational mechanics [21]) and of computational complexity [62].

Acknowledgments

A brief exposition of the local active information and local separable information were previously published as [63]. The authors thank Melanie Mitchell for helpful comments and suggestions regarding an early version of this manuscript. The CA diagrams were generated using enhancements to [64].

APPENDIX A: TOTAL INFORMATION COMPOSITION

Here, we demonstrate the mathematical arrangement of information storage and transfer for prediction of the next state at any given space-time point in the system.[75]

First, note that the *average* information required to predict the state at any spatiotemporal point is simply the single cell entropy H_X (Eq. (1)). We use the mutual information expansion of Eq. (5) to express this entropy in terms of the active information storage and entropy rate estimates[76]:

$$H_{X'} = I_{X';X^{(k)}} + H_{X'|X^{(k)}}. \quad (A1)$$

For convenience, we switch to local notation with the *local entropy* and *local temporal entropy rate* estimate represented as $h(n+1)$ and $h_\mu(n+1, k) = h(x_{i,n+1}|x_{i+r,n})$ respectively, and as $h(i, n+1)$ and $h_\mu(i, n+1, k)$ in local lattice notation:

$$h(n+1) = a(n+1, k) + h_\mu(n+1, k), \quad (A2)$$

$$h(i, n+1) = a(i, n+1, k) + h_\mu(i, n+1, k). \quad (A3)$$

Logically, we can restate this as: the information to predict a given cell site is the amount predictable from its past (the active memory) plus the remaining uncertainty after examining this memory.

We then consider the composition of this remaining uncertainty, tailoring our notation to CAs. In doing so, we alter Eq. (27) to represent a consecutive group of neighbors:

$$v_{i,n}^{s,f} = \{x_{i+q,n} | \forall q : s \leq q \leq f, q \neq 0\}. \quad (A4)$$

We systematically expand the entropy rate estimate term in Eq. (A3) by incrementally taking account of the contribution of each other information source. As a first step we identify the contribution from source $i+r$, the local apparent transfer entropy $t(i, -r, n+1, k)$ represented as a local conditional mutual information $m(x_{i+r,n}; x_{i,n+1} | x_{i,n}^{(k)})$ (see [55]):

$$\begin{aligned} h(i, n+1) = & a(i, n+1, k) + m(x_{i+r,n}; x_{i,n+1} | x_{i,n}^{(k)}) + \\ & h(x_{i,n+1} | x_{i+r,n}, x_{i,n}^{(k)}). \end{aligned} \quad (A5)$$

The systemic expansion is performed by using Eq. (7) iteratively on the rightmost conditional entropy term to produce:

$$\begin{aligned} h(i, n+1) = & a(i, n+1, k) + \\ & \sum_{j=-r, j \neq 0}^{+r} m(x_{i-j,n}; x_{i,n+1} | v_{i,n}^{-j+1,r}, x_{i,n}^{(k)}) + \\ & h(x_{i,n+1} | v_{i,n}^{-r,r}, x_{i,n}^{(k)}). \end{aligned} \quad (A6)$$

The sum of *incrementally conditioned mutual information terms* is a sum of transfer entropies from each element, incrementally conditioned on the previously considered sources in the neighborhood. The first such term in Eq. (A5) is an apparent transfer entropy, and the last term of the sum in Eq. (A6) is a complete transfer entropy. The order in which these *causal information contributors* are removed is arbitrary; the form of the result

will be the same. This sum is a *collective transfer entropy* from the set of all causal sources V to the destination: $I_{V;X'|X^{(k)}}$, i.e. the average information about the next state of the destination *jointly* added by the causal information sources that was not contained in the past of the destination. In local lattice notation, we have the *local collective transfer entropy* $t(i, n+1, k)$:

$$t(i, n+1, k) = \sum_{j=-r}^{+r} m(x_{i+j,n}; x_{i,n+1} | v_{i,n}^{-r,j-1}, x_{i,n}^{(k)}) \quad (A7)$$

$$= \log_2 \frac{p(x_{i,n+1} | x_{i,n}^{(k)}, v_{i,n}^{-r,r})}{p(x_{i,n+1} | x_{i,n}^{(k)})} \quad (A8)$$

As before, $t(i, n+1)$ represents the most correct form in the limit $k \rightarrow \infty$.

Also, we note that the final term in Eq. (A6) is the remaining *local intrinsic uncertainty* in the destination after its past and all causal information contributors have been considered; we label this as $u(i, n+1)$ (noting that it is independent of k). As such, Eq. (A6) logically displays the information to predict the destination at any

space-time point as the sum of the amount predictable from its past (active information storage), the amount then collectively predictable from its causal contributors (collective transfer entropy), and remaining intrinsic uncertainty:

$$h(i, n+1) = a(i, n+1, k) + t(i, n+1, k) + u(i, n+1). \quad (A9)$$

In a deterministic system, note that there is no remaining intrinsic uncertainty $u(i, n+1)$. As such, for deterministic systems (such as CAs) the temporal entropy rate $h_\mu(i, n+1, k)$ is equal to the collective transfer entropy $t(i, n+1, k)$.

As an example, the total information to predict the next state of any destination in an ECA can be represented as:

$$h(i, n+1) = a(i, n+1, k) + t(i, j=-1, n+1, k) + t^c(i, j=1, n+1, k), \quad (A10)$$

or vice-versa in $j = -1, 1$.

-
- [1] M. Mitchell, in *Non-Standard Computation*, edited by T. Gramss, S. Bornholdt, M. Gross, M. Mitchell, and T. Pellizzari (VCH Verlagsgesellschaft, Weinheim, 1998), pp. 95–140.
 - [2] M. Mitchell, in *20th Annual Conference of the Cognitive Science Society (Cogsci98)*, edited by M. A. Gernsbacher and S. J. Derry (Madison, Wisconsin, 1998), pp. 710–715.
 - [3] C. G. Langton, *Physica D* **42**, 12 (1990).
 - [4] S. Wolfram, *Physica D* **10**, 1 (1984).
 - [5] J. L. Casti, in *Beyond belief: randomness, prediction and explanation in science*, edited by J. L. Casti and A. Karlqvist (CRC Press, Boca Raton, 1991), pp. 280–327.
 - [6] M. Prokopenko, V. Gerasimov, and I. Tanev, in *Ninth International Conference on the Simulation of Adaptive Behavior (SAB'06)*, edited by S. Nolfi, G. Baldassarre, R. Calabretta, J. Hallam, D. Marocco, J.-A. Meyer, and D. Parisi (Springer Verlag, Rome, 2006), vol. 4095 of *Lecture Notes in Artificial Intelligence*, pp. 548–559.
 - [7] K. I. Goh and A. L. Barabási, *Europhysics Letters* **81**, 48002 (2008).
 - [8] R. Morgado, M. Cieřla, L. Longa, and F. A. Oliveira, *Europhysics Letters* **79**, 10002 (2007).
 - [9] J. A. Brown and J. A. Tuszynski, *Ferroelectrics* **220**, 141156 (1999).
 - [10] J. Pahle, A. K. Green, C. J. Dixon, and U. Kummer, *BMC Bioinformatics* **9**, 139 (2008).
 - [11] I. Couzin, R. James, D. Croft, and J. Krause, in *Fish Cognition and Behavior*, edited by B. C., K. Laland, and J. Krause (Blackwell Publishing, 2006), Fish and Aquatic Resources, pp. 166–185.
 - [12] A. S. Klyubin, D. Polani, and C. L. Nehaniv, in *8th European Conference on Artificial Life (ECAL 2005)*, edited by M. S. Capcarrere, A. A. Freitas, P. J. Bentley, C. G. Johnson, and J. Timmis (Springer Berlin / Heidelberg, Kent, UK, 2005), vol. 3630 of *Lecture Notes in Computer Science*, pp. 744–753.
 - [13] M. Lungarella and O. Sporns, *PLoS Computational Biology* **2**, e144 (2006).
 - [14] O. Kinouchi and M. Copelli, *Nature Physics* **2**, 348 (2006).
 - [15] J. J. Atick, *Network: Computation in Neural Systems* **3**, 213 (1992).
 - [16] M. A. Sánchez-Montañés and F. J. Corbacho, in *Proceedings of the International Conference on Artificial Neural Networks (ICANN 2002), Madrid, Spain*, edited by J. Dorronsoro (Springer-Verlag, Berlin/Heidelberg, 2002), vol. 2415 of *Lecture Notes in Computer Science*, pp. 141–141.
 - [17] T. Yamada and K. Aihara, in *IEEE Symposium on Emerging Technologies and Factory Automation (ETFA '94)* (IEEE, Tokyo, 1994), pp. 239–244.
 - [18] M. H. Jakubowski, K. Steiglitz, and R. Squier, *Phys. Rev. E* **56**, 7267 (1997).
 - [19] A. Adamatzky, ed., *Collision-Based Computing* (Springer-Verlag, Berlin, 2002).
 - [20] D. E. Edmundson and R. H. Enns, *Opt. Lett.* **18**, 1609 (1993).
 - [21] C. R. Shalizi, R. Haslinger, J.-B. Rouquier, K. L. Klinkner, and C. Moore, *Phys. Rev. E* **73**, 036104 (2006).
 - [22] J. E. Hanson and J. P. Crutchfield, *J. Stat. Phys.* **66**, 1415 (1992).
 - [23] J. Von Neumann, *Theory of self-reproducing automata* (University of Illinois Press, Urbana, 1966), edited and completed by Arthur W. Burks.
 - [24] M. Prokopenko, F. Boschetti, and A. J. Ryan (2008), Unpublished.
 - [25] D. J. MacKay, *Information Theory, Inference, and Learning Algorithms* (Cambridge University Press, Cam-

- bridge, 2003).
- [26] J. P. Crutchfield and D. P. Feldman, *Chaos* **13**, 25 (2003).
 - [27] P. Grassberger, *International Journal of Theoretical Physics* **25**, 907 (1986).
 - [28] W. Bialek, I. Nemenman, and N. Tishby, *Physica A* **302**, 89 (2001).
 - [29] S. Wolfram, *A New Kind of Science* (Wolfram Media, Champaign, IL, USA, 2002).
 - [30] L. Gray, *Notices of the American Mathematical Society* **50**, 200 (2003).
 - [31] A. Wuensche, *Complexity* **4**, 47 (1999).
 - [32] T. Helvik, K. Lindgren, and M. G. Nordahl, in *International Conference on Cellular Automata for Research and Industry*, edited by P. M. Soot, B. Chopard, and A. G. Hoekstra (Springer, Berlin/Heidelberg, Amsterdam, 2004), vol. 3305 of *Lecture Notes in Computer Science*, pp. 121–130.
 - [33] S. Wolfram, *Nature* **311**, 419 (1984).
 - [34] J. H. Conway, in *Winning ways for your mathematical plays*, edited by E. Berlekamp, J. H. Conway, and R. Guy (Academic Press, New York, 1982), vol. 2, ch. 25, pp. 927–962.
 - [35] M. Cook, *Comp. Sys.* **15**, 1 (2004).
 - [36] K. Lindgren and M. G. Nordahl, *Comp. Sys.* **4**, 299 (1990).
 - [37] M. H. Jakubowski, K. Steiglitz, and R. K. Squier, *Multiple-Valued Logic* **6**, 439 (2001).
 - [38] D. Eppstein, in *More Games of No Chance*, edited by R. J. Nowakowski (Cambridge Univ. Press, 2002), vol. 42 of *MSRI Publications*, pp. 433–453.
 - [39] M. Mitchell, J. P. Crutchfield, and P. T. Hraber, *Physica D* **75**, 361 (1994).
 - [40] M. Mitchell, J. P. Crutchfield, and R. Das, in *First International Conference on Evolutionary Computation and Its Applications*, edited by E. D. Goodman, W. Punch, and V. Uskov (Russian Academy of Sciences, Russia, Moscow, 1996).
 - [41] N. Boccara, J. Nasser, and M. Roger, *Phys. Rev. A* **44**, 866 (1991).
 - [42] G. J. Martinez, A. Adamatzky, and H. V. McIntosh, *Chaos, Solitons and Fractals* **28**, 100 (2006).
 - [43] W. Hordijk, C. R. Shalizi, and J. P. Crutchfield, *Physica D* **154**, 240 (2001).
 - [44] S. Wolfram, *Comm. Math. Phys.* **96**, 15 (1984).
 - [45] P. Grassberger, *J. Stat. Phys.* **45**, 27 (1986).
 - [46] H. V. McIntosh, *Linear Cellular Automata* (Universidad Autónoma de Puebla, Puebla, Mexico, 1990).
 - [47] R. Badii and A. Politi, *Phys. Rev. Lett.* **78**, 444 (1997).
 - [48] A. Lafusa and T. Bossomaier, in *The 2005 IEEE Congress on Evolutionary Computation* (IEEE Press, Edinburgh, 2005), vol. 1, pp. 844–849.
 - [49] H. Gutowitz and C. Domain, *Physica D* **103**, 155 (1997).
 - [50] P. Grassberger, in *Theory and Applications of Cellular Automata*, edited by S. Wolfram (World Scientific Publishing Co. Ltd., Singapore, 1986).
 - [51] K. Lindgren and M. G. Nordahl, *Comp. Sys.* **2**, 409 (1988).
 - [52] D. P. Feldman and J. P. Crutchfield, *Phys. Rev. E* **67**, 051104 (2003).
 - [53] A. S. Klyubin, D. Polani, and C. L. Nehaniv, in *Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems (ALife IX)*, Boston, edited by J. Pollack, M. Bedau, P. Husbands, T. Ikegami, and R. A. Watson (MIT Press, 2004), pp. 563–568.
 - [54] C. R. Shalizi, Ph.D. thesis, University of Wisconsin-Madison (2001).
 - [55] J. T. Lizier, M. Prokopenko, and A. Y. Zomaya, *Phys. Rev. E* **77**, 026110 (2008).
 - [56] J. E. Hanson and J. P. Crutchfield, *Physica D* **103**, 169 (1997).
 - [57] T. Schreiber, *Phys. Rev. Lett.* **85**, 461 (2000).
 - [58] *Oxford English Dictionary* (2008), accessed 8/5/2008, URL <http://www.oed.com/>.
 - [59] D. P. Feldman, C. S. McTague, and J. P. Crutchfield, *The organization of intrinsic computation: Complexity-entropy diagrams and the diversity of natural information processing* (2008), arXiv:0806.4789v1 [nlin.CD], URL <http://arxiv.org/abs/0806.4789>.
 - [60] J. T. Lizier, M. Prokopenko, and A. Y. Zomaya, in *Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems (ALife XI)*, Winchester, UK, edited by S. Bullock, J. Noble, R. Watson, and M. A. Bedau (MIT Press, Cambridge, MA, 2008), pp. 374–381.
 - [61] J. T. Lizier, M. Prokopenko, I. Tanev, and A. Y. Zomaya, in *Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems (ALife XI)*, Winchester, UK, edited by S. Bullock, J. Noble, R. Watson, and M. A. Bedau (MIT Press, Cambridge, MA, 2008), pp. 366–373.
 - [62] J. P. Crutchfield, *Physica D* **75**, 11 (1994).
 - [63] J. T. Lizier, M. Prokopenko, and A. Y. Zomaya, in *Advances in Artificial Life - 9th European Conference on Artificial Life (ECAL 2007)*, Lisbon, Portugal, edited by F. Almeida e Costa, L. M. Rocha, E. Costa, I. Harvey, and A. Coutinho (Springer, Berlin / Heidelberg, 2007), vol. 4648 of *Lecture Notes in Artificial Intelligence*, pp. 895–904.
 - [64] M. Wójtowicz, *Java celebration v.1.50* (2002), online Software, URL <http://psoup.math.wisc.edu/mcell/mjcell/mjcell.html>.
 - [65] A. Clauset, C. R. Shalizi, and M. E. J. Newman, *Power-law distributions in empirical data* (2007), arXiv:0706.1062v1, URL <http://arxiv.org/abs/0706.1062v1>.
 - [66] C. D. Manning and H. Schütze, *Foundations of Statistical Natural Language Processing* (MIT Press, Cambridge, MA, USA, 1999).
 - [67] N. Lüdtke, S. Panzeri, M. Brown, D. S. Broomhead, J. Knowles, M. A. Montemurro, and D. B. Kell, *Journal of The Royal Society Interface* (2007).
 - [68] H. Kantz and T. Schreiber, *Nonlinear Time Series Analysis* (Cambridge University Press, Cambridge, MA, 1997).
 - [69] $H_{\mu X}(k)$ here is equivalent to $h_{\mu}(k+1)$ in [26].
 - [70] In doing so, one should use proper quantitative techniques for establishing the presence or otherwise of power-laws, similar to [65] but tailored to study measurements of general functions rather than distributions only.
 - [71] This is as per Shalizi’s original formulation of the local excess entropy in [54], however our presentation is for a single time-series rather than the light-cone formulation used there. For a detailed description on why such average information-theoretical measures are the expectation value of local measures, and why the local measure can simply be taken as the log term within the averages measure, the reader is referred to the presentation of local transfer entropy in [55]. Also note that “lo-

cal” information-theoretic measures are known as “point-wise” measures elsewhere [66].

- [72] Calculation of $e(i, n, k)$ using local block entropies analogous to Eq. (12) would also require block entropies to be taken into the future to compute the same local information storage values. Without taking account of the dynamics into the future, we will not measure the information storage that *will* be used in the future of the process, but the information storage that is *likely* to be used in the future.
- [73] The separable information has parallels to the sum of “first order terms” (the apparent contribution of each source without considering interactions) in the total information of a destination in [67]. The local separable information however is distinguished in considering the contributions in the context of the destination’s past, and evaluating these on a local scale – both features are critical for an understanding of distributed computation.
- [74] The mutual information is measured over 200 000 sample points via kernel estimation of the underlying entropies [68], using a step kernel with precision $r = 0.3$ and maximum distance norms.
- [75] Similarly motivated is the decomposition of the total information for a destination presented in [67]. That approach however decomposes the total information in terms of a sum of mutual information terms of increasing size, rather than incrementally conditioned mutual information terms. Our presentation in the context of the destination’s history, and using local terms, is critical for an understanding of computation.
- [76] These equations are correct not only in the limit $k \rightarrow \infty$ but for estimates with any value of k .